arXiv:1110.2207v2 [cs.DS] 21 Nov 2011

Minimum Latency Submodular Cover

Sungjin Im[*] Viswanath Nagarajan [†] Ruben van der Zwaan [‡]

Abstract

We study the submodular ranking problem in the presence of metric costs. This is motivated by applications in broadcast scheduling and web-search ranking. The input to the *minimum latency submodular cover* (MLSC) problem consists of a metric (V, d) with source $r \in V$ and m monotone submodular functions $f_1, f_2, ..., f_m : 2^V \to [0, 1]$. The goal is to find a path originating at r that minimizes the total cover time of all functions; the cover time of function f_i is the smallest value t such that f_i has value one on the vertices visited within distance t along the path. We give a polynomial time $O(\log \frac{1}{\epsilon} \cdot \log^{2+\delta} |V|)$-approximation algorithm for this problem; here $\delta > 0$ is any constant, and ϵ is the minimum positive marginal increase of any function f_i. When the underlying metric is uniform, we obtain the *Ranking with Submodular Valuations* problem, for which an $O(\log \frac{1}{\epsilon})$-approximation algorithm was given by Azar and Gamzu [1]. In fact, our general result is enabled by a simpler analysis of [1].

We also consider the *latency covering Steiner tree* problem (LCST), which is an interesting special case of MLSC. Here, each function f_i is associated with a group $g_i \subseteq V$ and requirement k_i such that $f_i(S) = \min \left\{ \frac{|g_i \cap S|}{k_i}, 1 \right\}$. This is a natural generalization of the *latency group Steiner tree* [19, 8] and *generalized min-sum set cover* [2, 3, 25] problems. We obtain an $O(\log^2 |V|)$-approximation algorithm for LCST. This relies on a new and stronger LP relaxation for *covering Steiner tree*, that might be of independent interest. We note that any significant improvement over our results would also improve the best-known approximation for the well-studied *group Steiner tree* problem.

Finally, in a somewhat different direction, we consider the *stochastic submodular ranking* problem, and obtain an $O(\log 1/\epsilon)$-approximation algorithm, which is best possible. This problem models situations where web pages can be adaptively presented using user feedback. Our result also generalizes several previously studied stochastic problems, eg. [17, 26, 24].

1 Introduction

Ordering a set of elements so as to be simultaneously good for several valuations is an important issue in web-search ranking and broadcast scheduling. A formal model for this was introduced by Azar et al. [2] where they studied the *multiple intents re-ranking* problem (a.k.a. generalized min-sum set cover [3]). Subsequently, Azar and Gamzu [1] studied the *submodular ranking problem* where the valuations can be arbitrary (monotone) submodular functions.

In this paper, we extend both of these models to the setting of general metric switching costs. This allows us to handle additional issues such as:

---

[*] Department of Computer Science, University of Illinois, 201 N. Goodwin Ave., Urbana, IL 61801. `im3@illinois.edu`. Partially supported by NSF grant CCF-1016684. This work was partially done while the author was visiting IBM T. J. Watson Research Center.

[†] IBM T. J. Watson Research Center, Yorktown Heights, NY 10598. `viswanath@us.ibm.com`.

[‡] Department of Quantitative Economics, Maastricht University, The Netherlands. `r.vanderzwaan@maastrichtuniversity.nl`. This work was partially done while the author was visiting IBM T. J. Watson Research Center.

- *Data locality:* the time taken to read/transmit data j after data i is $d(i,j)$.
- *Context switching:* it takes $d(i,j)$ time for a user to parse data j when displayed after data i.

Thus we study the *minimum latency submodular cover* problem (MLSC) which is the metric version of submodular ranking [1], and the *latency covering Steiner tree* problem (LCST) which extends *generalized min-sum set cover* [2, 3]. We obtain poly-logarithmic approximation guarantees for both problems. We remark that due to a relation to the well-known *group Steiner tree* [16] problem, any significant improvement on our results would lead to a similar improvement for group Steiner tree.

In a somewhat different direction, we also study the *stochastic submodular ranking* problem, where the goal is to adaptively display stochastic elements so as to minimize the expected total cover time. We obtain an $O(\log \frac{1}{\epsilon})$-approximation algorithm for this problem, which is known to be best possible even in the deterministic setting [1].

1.1 Problem Definition

We let V denote the ground set of elements/vertices and $d : \binom{V}{2} \to \mathbb{R}_+$ the distance function. We assume that d is symmetric and satisfies triangle inequality. In addition there is a specified root vertex $r \in V$. There are m monotone submodular functions $f_1, \ldots, f_m : 2^V \to \mathbb{R}_+$ representing the valuations of different users. We assume WLOG by normalization that $f_i(V) = 1$ for all $i \in [m]$. Function f_i is said to be *covered* (or satisfied) by set $S \subseteq V$ if $f_i(S) = 1 = f_i(V)$. The *cover time* of function f_i in a path π is the length of the shortest prefix of π that has f_i value one, i.e.

$$\min\ t\ :\ f_i\left(\{v \in V : v \text{ appears within distance } t \text{ on } \pi\}\right) = 1.$$

The objective in the *minimum latency submodular cover problem* (MLSC) is to compute a path originating at r that minimizes the sum of cover times of all functions. We recover the *submodular ranking* problem [1] as the special case when metric d is uniform. A technical parameter that we use to measure performance (this was also used in [1]) is ϵ which is defined to be the smallest non-zero marginal increase of any function $\{f_i\}_{i=1}^m$.

As shown in Azar and Gamzu [1], the *submodular ranking* problem contains set-cover as a special case (even when $m = 1$). Similarly, we observe here that MLSC generalizes the *group Steiner tree* problem [16], where given metric (V, d) with root r and N groups of vertices $\{g_i \subseteq V\}_{i=1}^N$ the goal is find a minimum length tree containing r and at least one vertex from each of the N groups. Note that $f_1(S) = \frac{1}{N}\sum_{i=1}^N \min\{|g_i \cap S|, 1\}$ is a submodular function, with $f(S') = 1$ iff $S' \bigcap g_i \neq \emptyset$ for all $i \in [N]$. So MLSC on metric (V, d), root r and one function f_1 is equivalent to the corresponding *group Steiner tree problem* (up to factor 2 due to converting a tree into a path). Hence *group Steiner tree* is a special case of MLSC with $m = 1$ and $\epsilon = 1/N$. Similarly, *covering Steiner tree*, a generalization of *group Steiner tree* can be shown as a special case of MLSC.

A useful subroutine in our algorithm for MLSC is the *submodular orienteering problem* (SOP) [12] which is the following. The input is metric (V, d), root r, monotone submodular function $f : 2^V \to \mathbb{R}_+$ and length bound B. The goal is to find a path originating at r having length at most B that maximizes $f(S)$, where $S \subseteq V$ is the set of vertices visited in the path.

The *latency covering Steiner tree* problem (LCST) is a natural special case of MLSC, where each function f_i is associated with a group $g_i \subseteq V$ and requirement k_i such that $f_i(S) = \min\left\{\frac{|g_i \cap S|}{k_i},\, 1\right\}$. Note that in this case, $\epsilon = 1/\max_{i=1}^m k_i$. The uniform metric special case of LCST reduces to the *generalized* *min-sum set cover* problem [2, 3]. When $\max_{i=1}^{m} k_i = 1$ in LCST, we obtain the *latency group Steiner* problem, studied in Gupta et al. [19] (see also [27]) and Chakrabarty-Swamy [8].

Figure 1 shows the relations between the various problems discussed here.

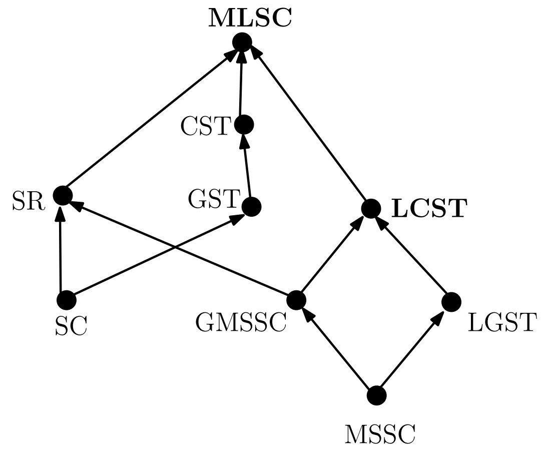

MLSC: minimum latency submodular cover

LCST : latency covering Steiner tree

CST : covering Steiner tree
GST : group Steiner tree
SR : submodular ranking
LGST : latency group Steiner tree
GMSSC : generalized min-sum set cover (multiple intents re-ranking)
SC : set cover
MSSC : min sum set cover

Figure 1: The various problems and their relations. An arrow from X to Y means X is a special case of Y.

Stochastic submodular ranking. Here we are given a set V of stochastic elements, each having an independent distribution over certain domain Δ. The submodular functions are also defined on groundset Δ, i.e. $f_1, ..., f_m : 2^{\Delta} \to [0, 1]$. In addition, each element $i \in V$ has a deterministic cost/time ℓ_i to be scheduled.[1] The realization (from Δ) of any element is known immediately after scheduling/displaying it. The goal is to find an adaptive ordering of V that minimizes the total expected cover time. Since elements are stochastic, it is possible that a function is never covered: in such cases we just fix the cover time to be $\sum_{i \in V} \ell_i$ (which is the total duration of any schedule).

We will be concerned with *adaptive* algorithms. Such an algorithm is allowed to decide the next element to display based on the instantiations of the previously scheduled elements. This models the setting where the algorithm can benefit from user feedback. Our stochastic problem clearly generalizes the *stochastic set cover* studied in [17]. Interestingly, it also captures some variants of *stochastic set cover* that have applications in processing multiple queries with probabilistic information [26, 24]. Various applications of *stochastic submodular ranking* are discussed in more detail in Section 7.

1.2 Our Results and Techniques

We start with the general MLSC problem, for which we obtain a poly-logarithmic approximation.

Theorem 1.1. *For any constant $\delta > 0$, there is an $O(\log \frac{1}{\epsilon} \cdot \log^{2+\delta} |V|)$-approximation algorithm for the minimum latency submodular cover problem.*

---

[1] We remark that one could also consider a stochastic generalization of MLSC, where the costs form a metric; however our algorithms currently do not apply in this generality.

Note that in the special case of *group Steiner tree*, this result is only factor $O(\log^{\delta} |V|)$ larger than its best known approximation ratio of $O(\log N\ \log^2 |V|)$, due to Garg et al. [16]. In fact our algorithm for MLSC is essentially a reduction to *submodular orienteering* (SOP). Specifically, we show that an (α, β)-bicriteria approximation algorithm[2] for SOP can be used to obtain an $O(\alpha\,\beta \cdot \log \frac{1}{\epsilon})$-approximation algorithm for MLSC. To obtain Theorem 1.1 we use an $\big(O(1),\, O(\log^{2+\delta} |V|)\big)$-bicriteria approximation for SOP that follows from [6, 11]. Alternatively, we can use the *quasi-polynomial* time true $O(\log |V|)$-approximation for SOP by Chekuri and Pal [12] within our algorithm to obtain:

Theorem 1.2. *There is a quasi-polynomial time* $O(\log \frac{1}{\epsilon} \cdot \log |V|)$*-approximation algorithm for the minimum latency submodular cover problem.*

Our algorithm for MLSC is an extension of the elegant "adaptive residual updates scheme" of Azar and Gamzu [1] for *submodular ranking* (i.e. uniform metric MLSC). As shown in [1], an interesting aspect of this problem is that the natural greedy algorithm, based on absolute contribution of elements, performs very poorly. Instead they used a modified greedy algorithm that selects one element at a time according to residual coverage. In the MLSC setting of general metrics, our algorithm uses a similar *residual coverage function* to repeatedly augment the solution. However our augmentations are paths of geometrically increasing lengths, instead of just one element. A crucial point in our algorithm is that the residual coverage functions are always submodular, and hence we can use *submodular orienteering* (SOP) in the augmentation step. We remark that the approach of covering the maximum number of functions within geometrically increasing lengths fails because the residual coverage function here is non-submodular; in fact as noted in [3] this subproblem contains the difficult dense-k-subgraph problem (even for generalized min-sum set cover with requirement two). We also note that the choice of our (submodular) residual coverage function ultimately draws on the submodular ranking algorithm [1].

The analysis in [1] was based on viewing the optimal and approximate solutions as histograms. This approach was first used in this line of work by Feige et al. [15] for the *min-sum set cover* problem (see also [4]). This was also the main framework of analysis in [2] for *generalized min-sum set cover* and then for *submodular ranking* [1]. However, these proofs have been increasingly difficult as the problem in consideration adds more generality. Instead we follow a different and more direct approach that is similar to analysis for the *minimum latency TSP*, see eg. [10, 13]. In Section 3 we first give a simpler proof of the submodular ranking result [1], which also implies the result in Azar et al. [2]. Then we present our algorithm for MLSC in Section 4, which yields Theorems 1.1 and 1.2.

Our second main result is for the LCST problem:

Theorem 1.3. *There is an* $O(\log^2 |V|)$*-approximation algorithm for latency covering Steiner tree.*

Note that our result for the more general MLSC problem implies directly for LCST (i) polynomial time $O(\log k_{max} \cdot \log^{2+\delta} |V|)$-approximation and (ii) quasi-polynomial time $O(\log k_{max} \cdot \log |V|)$-approximation. Here k_{max} denotes the maximum requirement of any group in LCST. Theorem 1.3 improves on these bounds by providing a double logarithmic approximation in polynomial time. As shown in [27], any improvement over Theorem 1.3 even in the $k_{max} = 1$ special case (i.e. *latency group Steiner tree*) would yield an improved approximation ratio for *group Steiner tree*, which is a long-standing open question.

Recall that LCST on a uniform metric is the *generalized min-sum set cover* problem, for which a constant factor approximation is known due to Bansal et al. [3]; the previous greedy-type algorithm [2] had a logarithmic ratio. A natural attempt at solving LCST is to extend the algorithm in [3] to general metrics. However, this is not straightforward since the [3] algorithm is LP-based, and there is no good LP relaxation known for the underlying *covering Steiner tree* problem. Although the $O(\log N \ \log^2|V|)$-approximation algorithm for *covering Steiner tree* [20, 23] uses an LP relaxation, that LP is known to have a large integrality gap [23]. The algorithms [20, 23] get around this integrality gap by iteratively (and adaptively) solving a sequence of such LPs; thus the final solution cost cannot be compared against a fixed LP relaxation. In context of the LCST problem it is unclear how to use this approach to get an $O(\log^2|V|)$-approximation.[3] We overcome this issue by proposing a new LP relaxation for covering Steiner tree, that has a small (polylogarithmic) integrality gap. The new LP makes use of *Knapsack Cover* (KC) type inequalities [7]. We note that KC-inequalities (for set multicover) were also used in [3]. Still, it is non-trivial to formulate these stronger inequalities for *covering Steiner tree* in a way that they are useful in rounding. In particular, the separation oracle for the resulting KC-inequalities is more complicated than for *set multicover* (which was the knapsack problem). Once we have the stronger LP relaxation for *covering Steiner tree*, we can use it in an LP-based algorithm for LCST using fairly standard techniques. These results are presented in Section 5 and Section 6, respectively.

---

[2] Note that an (α, β) bicriteria guarantee for SOP means that on every instance, the algorithm produces a path of length at most β times the bound and profit at least $\frac{1}{\alpha}$ times the optimum (that strictly satisfies the length bound).

Our final result is on the *stochastic submodular ranking* problem. As shown in [17, 18], even special cases of this problem have polynomially large adaptivity gap (ratio between the optimal non-adaptive and adaptive solutions). This motivates us to focus on adaptive algorithms, and we obtain the following result in Section 7.

Theorem 1.4. *There is an adaptive $O(\log \frac{1}{\epsilon})$-approximation algorithm for the stochastic submodular ranking problem.*

In particular, we show that the most natural stochastic extension of the algorithm from [1] achieves the approximation factor. We remark that the analysis in [1] for deterministic submodular ranking assumed unit costs, whereas Theorem 1.4 holds for the stochastic setting even with non-uniform costs $\{\ell_i\}$.

As mentioned before, our results generalize the results in [17, 26, 24] which study (some variants of) stochastic set cover problems. Our analysis is arguably simper and more transparent than [24], which gave the first tight analysis for these problems. We note that the charging scheme used in [24] makes use of "dual prices" and it does not seem directly applicable to general submodular functions.

1.3 Previous Work

The *min sum set cover* problem was studied in Feige et al. [15] where they proved an approximation ratio of 4 for the natural greedy algorithm, and also showed a matching hardness of approximation. The algorithmic result was also implicit in Bar-Noy et al. [4]. Azar et al. [2] introduced the *generalized min-sum set cover* problem and gave a modified greedy algorithm having an $O(\log k_{max})$-approximation ratio, where k_{max} is the maximum requirement of any set. Their analysis built on the histogram based approach from [15]. Later Bansal et al. [3] gave an improved $O(1)$-approximation using LP and knapsack cover inequalities [7]. Recently, Skutella and Williamson [25] improved the approximation ratio to 28 using alpha-point rounding methods from scheduling.

The *submodular ranking* problem was introduced by Azar and Gamzu [1] where they obtained a tight $O(\log \frac{1}{\epsilon})$-approximation algorithm using a nice residual updates scheme. The analysis was histogram based and fairly involved. In contrast, we use a more direct approach that also extends to the general metric setting.

---

[3] The adaptive residual coverage algorithm for MLSC is somewhat similar to adaptively solving LPs for covering Steiner tree; but that approach gives us only a weaker bound for LCST.

The first poly-logarithmic approximation for group Steiner tree was $O(\log N\ \log^2 |V|)$ due to Garg et al. [16] which is still the best known bound. Chekuri et al. [11] gave a combinatorial recursive greedy algorithm that achieved a slightly weaker approximation ratio (the [16] result was LP-based). This combinatorial approach was extended in Calinescu and Zelikovsky [6] to the problem of covering any submodular function in a metric. We use this algorithm in the submodular orienteering (SOP) subroutine for our MLSC result. For SOP a true $O(\log |V|)$-approximation is known due to Chekuri and Pal [12], but running in *quasi-polynomial* time. We note that an $\Omega(\log^{2-\delta} |V|)$ hardness of approximation is known for *group Steiner tree* (even on tree metrics) due to Halperin and Krauthgamer [21].

The *covering Steiner tree* problem was introduced in Konjevod et al. [23], which can be viewed as the multicover version of *group Steiner tree*. They gave an $O(\log(Nk_{max})\ \log^2 |V|)$-approximation using an LP-relaxation. However the LP used in [23] has a large $\Omega(k_{max})$ integrality gap; they got around this issue by iteratively solving a suitable sequence of LPs. They also extended the randomized rounding analysis from [16] to this context. Later, Gupta and Srinivasan [20] improved the approximation bound to $O(\log N\ \log^2 |V|)$, removing the dependence on the covering requirements. This algorithm was also based on solving a similar sequence of LPs; but the improvement was due to a combination of threshold-rounding and randomized rounding. In this paper, we give a stronger LP relaxation for *covering Steiner* based on KC-inequalities, that has an $O(\log N\ \log^2 |V|)$ integrality gap. This might be of some independent interest. This new LP is also crucial for our $O(\log^2 |V|)$-approximation algorithm for LCST.

The *latency group Steiner* problem was studied in [19, 27] and also [8], in different contexts. These papers imply an $O(\log^2 |V|)$-approximation algorithm for this problem.

The *adaptive set cover* problem introduced in [17] is clearly a special case of stochastic SR that we consider. Goemans and Vondrak [17] showed a large adaptivity gap for set cover, and a logarithmic approximation for a variant with multiplicities. A related problem in context of fast query evaluation was studied in Munagala et al. [26], where the authors gave a triple logarithmic approximation. This bound was improved to the optimal logarithmic ratio by Liu et al. [24]; this result was also applicable to adaptive set cover. Another related paper is by Golovin and Krause [18], where they defined a general property "adaptive submodularity" and showed nearly optimal approximation guarantees for several objectives (max coverage, min-cost cover and min-sum cover). The most relevant result in [18] to the stochastic SR problem is the 4-approximation for stochastic *min sum set cover*. This approach required a *fixed* submodular function f such that the objective is $\mathbb{E}\left[\sum_{t\geq 0} f(\overline{V}) - f(\overline{\pi}_t)\right]$ where $\overline{\pi}_t$ is the instantiation of elements displayed within time t and $\overline{V}$ denotes the instantiation of all elements. However, as mentioned earlier this is not the case even for generalized min-sum set cover with requirements two.

2 Preliminaries

Let V denote a ground set of elements. A function $f : 2^V \to \mathbb{R}$ is said to be submodular if it satisfies the property that for any $A, B \subseteq V$, $f(A)+f(B) \geq f(A\cup B)+f(A\cap B)$. It is well-known that this is equivalent to the following property: For any $A \subseteq B$ and any $e \in V \setminus B$, $f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B)$.

Many interesting objectives can be modeled as submodular functions. For example, in economics, the latter property of a submodular fuction can be viewed essentially as the law of diminishing marginal utility. Also submodular functions appear often in combinatorial optimization. Such examples include set unions, cut functions, matroid rank functions etc. We refer the reader to [28] for an extensive overview of submodular functions and their applications.

From the definition of submodularity, the following useful inequality can be derived.

Proposition 2.1. *For any $A, B \subseteq V$ and submodular function f, it follows that*

$$\sum_{e \in A} \Big(f(B \cup \{e\}) - f(B) \Big) \geq f(B \cup A) - f(B)$$

A submodular function f is said to be monotone if for any $A \subseteq B$, $f(A) \leq f(B)$. Since function f is defined over all possible subsets of V, it is a standard assumption that a value oracle for f is available. The oracle takes as input a subset $S \subseteq V$ and returns the function value $f(S)$ in constant time.

Notation: We will use ALG to denote the cost of the algorithm we consider. For notational simplicity, we will let OPT denote the optimal solution itself or the optimal cost depending on the context. For any integer $k \geq 1$ we let $[k] := \{1, 2, \ldots, k\}$. For notational convenience, we will allow the same notation to denote a tree or the set of edges in the tree.

3 Simpler Analysis of the Submodular Ranking Problem

In this section, we revisit the *submodular ranking* problem [1]. Recall that the input consists of a ground set $V := [n]$ of elements and monotone submodular functions $f_1, f_2, ...f_m : 2^{[n]} \rightarrow [0, 1]$ with $f_i(V) = 1$, $\forall i \in [m]$. The goal is to find a complete linear ordering of the elements that minimizes the total cover time of all functions. The cover time $\mathsf{cov}(f_i)$ of f_i is defined as the smallest index t such that the function f_i has value 1 on the first t elements in the ordering. We also say that an element e is scheduled at time t if it is the t^{th} element in the ordering. It is assumed that each function f_i satisfies the following property: for any $S \supseteq S'$, if $f_i(S) - f_i(S') > 0$ then it must be the case that $f_i(S) - f_i(S') \geq \epsilon$, where $\epsilon > 0$ is a constant that is uniform for all functions f_i. This is a useful parameter in describing the performance guarantee.

Azar and Gamzu [1] gave a modified greedy-style algorithm with an approximation factor of $O(\log \frac{1}{\epsilon})$ for *submodular ranking*. Their analysis was histogram-based and fairly involved. In this section, we give an alternate proof of their result. We believe that our analysis is easier and perhaps more intuitive. Our analysis also extends to the more general MLSC problem which we will study in the next section. The algorithm ALG-AG from [1] is given below. In the output, $\pi(t)$ denotes the element that appears in the tth time slot.

Algorithm 1 ALG-AG

```
INPUT: Ground set $[n]$; monotone submodular functions $f_i : 2^{[n]} \rightarrow [0, 1], i \in [m]$
$S \leftarrow \emptyset$
for $t = 1$ to $n$ do
    Let $f^S(e) := \sum_{i \in [m], f_i(S) < 1} \frac{f_i(S \cup \{e\}) - f_i(S)}{1 - f_i(S)}$
    $e = \arg\max_{e \in [n] \setminus S} f^S(e)$
    $S \leftarrow S \bigcup \{e\}$
    $\pi(t) \leftarrow e$
end for
OUTPUT: A linear ordering $\langle \pi(1), \pi(2), \ldots, \pi(n) \rangle$ of $[n]$.
```

Theorem 3.1 ([1])**.** *The algorithm ALG-AG is an $O(\ln(\frac{1}{\epsilon}))$-approximation for submodular ranking.*

Let $\alpha := 1 + \ln(\frac{1}{\epsilon})$. To simplify notation, WLOG, we assume that α is an integer. Let $R(t)$ denote the set of functions that are not satisfied by ALG-AG earlier than time t; $R(t)$ includes the functions that are satisfied exactly at time t. For notational convenience, we will use $i \in R(t)$ interchangeably with

$f_i \in R(t)$. Analogous set $R^*(t)$ is defined for the optimal solution. Note that $\mathsf{ALG} = \sum_{t \geq 1} |R(t)|$ and $\mathsf{OPT} := \sum_{t \geq 1} |R^*(t)|$. We will be interested in the number of unsatisfied functions at times $\{8\alpha 2^j \ : \ j \in \mathbb{Z}_+\}$ by ALG-AG and the number of unsatisfied functions at times $\{2^j \ : \ j \in \mathbb{Z}_+\}$ by the optimal solution. Let $R_j := R(8\alpha 2^j)$ and $R^*_j = R^*(2^j)$. It is important to note that R_j and R^*_j are concerned with different times. For notational simplicity, we let $R_{-1} := \emptyset$.

We will show the following key lemma. Roughly speaking, it says that the number of unsatisfied functions by ALG-AG diminishes quickly unless it is comparable to the number of unsatisfied functions in OPT.

Lemma 3.2. *For any* $j \geq 0$, *we have* $|R_j| \leq \frac{1}{4}|R_{j-1}| + |R^*_j|$.

Proof of Theorem 3.1. Assuming Lemma 3.2, we proceed as follows.

$$\begin{aligned}
\mathsf{ALG} &= \sum_{j \geq 0} \sum_{8\alpha 2^j \leq t < 8\alpha 2^{j+1}} |R(t)| \quad + \sum_{1 \leq t < 8\alpha} |R(t)| \\
&\leq \sum_{j \geq 0} 8\alpha(2^{j+1} - 2^j)|R_j| \quad + \quad 8\alpha \mathsf{OPT} \\
&\qquad \text{[Since } |R(t)| \text{ is non-increasing, and for any } t \geq 1, |R(t)| \leq m \leq \mathsf{OPT}] \\
&= 8\alpha \sum_{j \geq 0} 2^{j+1} \left(|R_j| - \frac{1}{4}|R_{j-1}| \right) \quad + \quad 8\alpha \mathsf{OPT} \\
&\leq 8\alpha \sum_{j \geq 0} 2^{j+1} |R^*_j| \quad + \quad 8\alpha \mathsf{OPT} \qquad \text{[By Lemma 3.2]} \\
&\leq 32\alpha \sum_{j \geq 1} \sum_{2^{j-1} \leq t < 2^j} |R^*_j| \ + \ 16\alpha |R^*_0| \quad + \quad 8\alpha \mathsf{OPT} \\
&\leq 32\alpha \mathsf{OPT} \ + \ 24\alpha \mathsf{OPT} \qquad \text{[Since } |R^*(t)| \text{ is non-increasing]}
\end{aligned}$$

Thus we obtain $\mathsf{ALG} \leq 56\alpha \, \mathsf{OPT}$, which proves Theorem 3.1.

Henceforth, we will focus on proving Lemma 3.2. Consider any time step $t \in [8\alpha 2^{j-1}, 8\alpha 2^j)$ for an integer $j \geq 1$; when $j = 0$ the lemma trivially holds. Let S_{t-1} denote the set of elements that ALG-AG displays before time t and let e_t denote the element that ALG-AG schedules exactly at time t. Let E_j denote the set of elements that ALG-AG schedules until time $8\alpha 2^j$. Let E^*_j denote the set of elements that OPT schedules until time 2^j. Recall that ALG-AG picks e_t as an element e that maximizes

$$f^{S_{t-1}}(e) := \sum_{i \in [m] : f_i(S_{t-1}) < 1} \frac{f_i(S_{t-1} \cup \{e\}) - f_i(S_{t-1})}{1 - f_i(S_{t-1})}$$

This leads us to the following proposition.

Proposition 3.3. *Consider any* $j \geq 1$ *and any time step* $t \in [8\alpha 2^{j-1}, 8\alpha 2^j)$. *Then for any* $e \in E^*_j$ *, we have* $f^{S_{t-1}}(e_t) \geq f^{S_{t-1}}(e)$.

Proof. Since ALG-AG has chosen to schedule element e_t over all elements $e \in E^*_j \setminus S_{t-1}$, we know that the claimed inequality holds for any $e \in E^*_j \setminus S_{t-1}$. Further, the inequality holds for any element e in S_{t-1}, since $f^{S_{t-1}}(e) = 0$ for such an element e. □

By taking an average over all elements in E_j^*, we derive

$$\begin{aligned} f^{S_{t-1}}(e_t) &\geq \frac{1}{|E_j^*|}\sum_{e\in E_j^*} f^{S_{t-1}}(e) \\ &\geq \frac{1}{|E_j^*|}\sum_{e\in E_j^*}\sum_{i\in R_j\setminus R_j^*}\frac{f_i(S_{t-1}\cup\{e\})-f_i(S_{t-1})}{1-f_i(S_{t-1})} \end{aligned} \tag{1}$$

Observe that in the second inequality, we can drop the constraint $f_i(S_{t-1}) < 1$ in the subscript, since S_{t-1} does not cover f_i for any $i \in R_j$: by definition, f_i is not covered before time $8\alpha 2^j$ and $t < 8\alpha 2^j$. Due to submodularity of each function f_i (more precisely by applying Proposition 2.1 with $A := E_j^*$ and $B := S_{t-1}$), we have that

$$\begin{aligned} (1) &\geq \frac{1}{|E_j^*|}\sum_{i\in R_j\setminus R_j^*}\frac{f_i(S_{t-1}\cup E_j^*)-f_i(S_{t-1})}{1-f_i(S_{t-1})} \\ &= \frac{1}{|E_j^*|}\sum_{i\in R_j\setminus R_j^*} 1 \geq \frac{|R_j|-|R_j^*|}{|E_j^*|} \end{aligned}$$

The equality is due to the fact that for any $i \notin R_j^*$, $f_i(E_j^*) = 1$ and each function f_i is monotone. Hence:

$$\sum_{8\alpha\cdot 2^{j-1}\leq t<8\alpha\cdot 2^j} f^{S_{t-1}}(e_t) \;\geq\; \frac{8\alpha(2^j-2^{j-1})}{|E_j^*|}(|R_j|-|R_j^*|) \;=\; 4\alpha(|R_j|-|R_j^*|), \tag{2}$$

where we used $|E_j^*| = 2^j$. We now upper bound the left-hand-side of (2). To this end, we need the following claim from [1].

Claim 3.4 (Claim 2.3 in [1])**.** *Given a monotone function* $f : 2^{[n]} \to [0,1]$ *and sets* $\emptyset = S_0 \subseteq S_1 \subseteq \cdots \subseteq S_\ell \subseteq [n]$*, we have (using the convention* $0/0 = 0$*)*

$$\sum_{k=1}^{\ell}\frac{f(S_k)-f(S_{k-1})}{1-f(S_{k-1})} \;\leq\; 1+\ln\frac{1}{\delta}.$$

Here $\delta > 0$ *is such that for any* $A \subseteq B$*, if* $f(B)-f(A) > 0$ *then* $f(B)-f(A) \geq \delta$*.*

Proof. We give a proof for completeness. Order the values in the set $\{f(S_k) \mid 0 \leq k \leq \ell\} \setminus \{1\}$ in increasing order to obtain $\beta_0 < \beta_1 < \ldots < \beta_H$. By the assumption, we have $\beta_0 = 0$, $1-\beta_H \geq \delta$ and $\beta_h - \beta_{h-1} \geq \delta$, $\forall h \in [H]$. Note that the desired inequality is reduced to:

$$\sum_{h=1}^{H}\frac{\beta_h-\beta_{h-1}}{1-\beta_{h-1}} \quad\leq\quad \ln\frac{1}{\delta}$$

If $f(S_\ell) = 1$ then the summation we want to bound has an additional term of $\frac{1-\beta_H}{1-\beta_H} = 1$.

Knowing that the function $x \to \frac{1}{1-x}$ is increasing for $x \in [0,1)$, we derive

$$\begin{aligned}\sum_{h=1}^{H} \frac{\beta_h - \beta_{h-1}}{1-\beta_{h-1}} &= \sum_{h=1}^{H} \int_{x=\beta_{h-1}}^{\beta_h} \frac{1}{1-\beta_{h-1}}\, \mathsf{d}x \le \sum_{h=1}^{H} \int_{x=\beta_{h-1}}^{\beta_h} \frac{1}{1-x}\, \mathsf{d}x = \int_{x=0}^{\beta_H} \frac{1}{1-x}\, \mathsf{d}x \\ &= -\ln(1-\beta_H) \quad \le \quad \ln\frac{1}{\delta}\end{aligned}$$

This proves the claim. ☐

Note that any function f_i not in R_{j-1} does not contribute to the left-hand-side of (2). This is because such a function f_i was already covered before time $8\alpha\, 2^{j-1} \le t$. Further, knowing by Claim 3.4 that each function $f_i \in R_{j-1}$ can add at most $\alpha := 1 + \ln\frac{1}{\epsilon}$, we can upper bound the left-hand-side of (2) by $\alpha|R_{j-1}|$. More formally,

$$\begin{aligned}\sum_{8\alpha\cdot 2^{j-1} < t \le 8\alpha\cdot 2^j} f^{S_{t-1}}(e_t) &= \sum_{8\alpha\cdot 2^{j-1} < t \le 8\alpha\cdot 2^j} \; \sum_{i\in R_{j-1}: f_i(S_{t-1})<1} \frac{f_i(S_{t-1}\cup\{e_t\}) - f_i(S_{t-1})}{1-f_i(S_{t-1})} \\ &\le \sum_{i\in R_{j-1}} \; \sum_{t\ge 1: f_i(S_{t-1})<1} \frac{f_i(S_{t-1}\cup\{e_t\}) - f_i(S_{t-1})}{1-f_i(S_{t-1})} \le \alpha|R_{j-1}| \qquad (3)\end{aligned}$$

From (2) and (3) we obtain

$$4\alpha(|R_j| - |R_j^*|) \le \alpha|R_{j-1}|$$

which simplifies to the desired inequality in Lemma 3.2.

4 Minimum Latency Submodular Cover

Recall that in the *minimum latency submodular cover* problem (MLSC), we are given a metric (V,d) with root $r \in V$ and m monotone submodular functions $f_1, f_2, ..., f_m : 2^V \to [0,1]$. Each function f_i satisfies the following properties: (1) $f_i(V) = 1$ and (2) for any $S' \supset S$, if $f_i(S') - f_i(S) > 0$, then it must be the case that $f_i(S') - f_i(S) \ge \epsilon$, where $\epsilon > 0$ is a value that is uniform for all functions. WLOG, we assume that $d(\cdot,\cdot)$ is an integer. The goal is to find a path starting at r that minimizes the total cover time of all functions.

As mentioned earlier, our algorithm for MLSC uses as sub-routine an algorithm for the Submodular Orienteering Problem (SOP). Here given metric (V,d), root r, monotone submodular function $g : 2^V \to \mathbb{R}_+$ and bound B, the goal is to compute a path P originating at r having length at most B that maximizes $g(V(P))$ where $V(P)$ denotes the set of vertices covered by P. We assume a $(\rho,\, \sigma)$-bicriteria approximation algorithm ALG-SOP for SOP, i.e. on any SOP instance returns a path P of length at most $\sigma \cdot B$ and $g(V(P)) \ge \mathsf{OPT}/\rho$. We recall the following known results on SOP.

Theorem 4.1 ([6])**.** *For any constant $\delta > 0$ there is a polynomial time $(O(1),\, O(\log^{2+\delta}|V|))$ bicriteria approximation algorithm for SOP.*

Theorem 4.2 ([12])**.** *There is a quasi-polynomial time $O(\log|V|)$ approximation algorithm for SOP.*

We now describe our algorithm ALG-LSCP for MLSC that uses the (ρ, σ) bicriteria approximation algorithm ALG-SOP. Here $\alpha = 1 + \ln \frac{1}{\epsilon}$. Note the difference from the submodular ranking algorithm [1]: here each augmentation is a path possibly covering several vertices. Despite the similarity of ALG-LSCP to the min-latency TSP type algorithms [10, 13] an important difference is that we *do not* try to directly maximize the number of covered functions in each augmentation: as noted before this subproblem is at least as hard as dense-k-subgraph, for which the best approximation ratio known is only polynomial [5]. Instead we maximize in each step some proxy residual coverage function f^S that suffices to eventually cover all functions quickly. This function is a natural extension of the single-element coverage values used in ALG-AG [1]. It is important to note that in Line (4), $f^S(\cdot)$ is defined adaptively based on the current set S of visited vertices in each iteration. Moreover, since each function f_i monotone submodular, so is f^S for any $S \subseteq V$. In Line (6), $\pi \cdot P$ implies the concatenation of π and P.

Algorithm 2 ALG-LSCP

```
INPUT: $(V, d), r \in V; \{f_i : 2^V \to [0,1]\}_{i=1}^{m}$.
1: $S \leftarrow \emptyset$, $\pi \leftarrow \emptyset$.
2: for $k = 0, 1, 2, ...$ do
3:    Repeat the following $4\alpha\rho$ times.
4:        Define submodular function $f^S(T) := \sum_{i \in [m], f_i(S) < 1} \frac{f_i(S \cup T) - f_i(S)}{1 - f_i(S)}$, for all $T \subseteq V$.
5:        Use ALG-SOP to find a path $P$ of length at most $\sigma \cdot 2^k$ starting from $r$ that maximizes $f^S(V(P))$
          (within a factor of $\rho$) where $V(P)$ is the set of nodes visited by $P$.
6:        $S \leftarrow S \cup V(P)$ and $\pi \leftarrow \pi \cdot P$.
7: end for
OUTPUT: Output solution $\pi$.
```

We prove the following theorem, which implies Theorem 1.1 and 1.2.

Theorem 4.3. *ALG-LSCP is an $O(\alpha\rho\sigma)$-approximation algorithm for MLSC.*

We now analyze ALG-LSCP. We will say that the algorithm is in the jth phase, when the variable k of the for loop in Line (2) has value j. Observe that the final solution visits all vertices that are added in the jth phase within time $16\alpha\rho 2^j$. This can be easily shown as follows: the final solution is a concatenation of the paths that were found in Line (5). Since all these paths are stitched at the root r, the length of π at the end of phase j is at most $\sum_{k=1}^{j} 2 \cdot 4\alpha\rho \cdot \sigma 2^k \leq 16\alpha\rho\sigma \cdot 2^j$. The following proposition easily follows.

Proposition 4.4. *Any vertex v added to S in the jth phase is visited by π within $16\alpha\rho\sigma \cdot 2^j$.*

Let $R(t)$ denote the set of (indices of) the functions that are not covered by ALG-LSCP earlier than time t; $R(t)$ includes the functions that are covered exactly at time t as well. We will interchangeably use $i \in R(t)$ from $f_i \in R(t)$ for notational simplicity. Let $R_j := R_j(16\alpha\rho\sigma\, 2^j)$. Similarly, we let $R^*(t)$ denote the set of functions that are not covered by OPT earlier than time t and let $R^*_j = R^*(2^j)$. For notational convenience, we let $R_{-1} := \emptyset$.

We will show the following key lemma. It shows that the number of uncovered functions by ALG-LSCP must decrease fast as j grows, unless the corresponding number by the optimal solution is comparable.

Lemma 4.5. *Consider any $j \geq 0$. Then we have $|R_j| \leq \frac{1}{4}|R_{j-1}| + |R^*_j|$.*

Assuming the above lemma holds, we can prove Theorem 4.3 exactly as we proved Theorem 3.1 in the previous section using Lemma 3.2. Thus it suffices to prove Lemma 4.5, which we do next.

The lemma trivially holds when $j = 0$, hence consider any fixed phase $j \geq 1$. Let S_0 denote the set of vertices that were add to S up to the end of phase $j-1$. Let $H = 4\alpha\rho$ and $T_1, T_2, ..., T_H$ be the sets of vertices that were added in Line (6) in the jth phase. Let $S_h = S_0 \cup T_1 \cup T_2 \cup ... \cup T_h, \forall 1 \leq h \leq H$. We will prove Lemma 4.5 by lower and upper bounding the quantity

$$\Delta_j \ := \ \sum_{h=1}^{H} f^{S_{h-1}}(T_h) \ = \ \sum_{h=1}^{H} \ \sum_{i \in [m]: f_i(S_{h-1}) < 1} \frac{f_i(S_h) - f_i(S_{h-1})}{1 - f_i(S_{h-1})}$$

We first lower bound Δ_j. Let T^* denote the set of vertices that OPT visited within time 2^j. Observe that in Line (5), ALG-LSCP could have visited all nodes in T^* by choosing P as the prefix of length 2^j of OPT. Via the approximation guarantee of ALG-SOP, we obtain

Proposition 4.6. *For any $h \in [H]$ and $i \in R_j$, we have $f^{S_{h-1}}(S_h) \geq \frac{1}{\rho} \cdot f^{S_{h-1}}(T^*)$.*

For the remaining proof, we only give a sketch, since it is very similar to the proof of Lemma 3.2. We restrict our concern to the functions in $R_j \setminus R_j^*$. Observe that for any $i \in R_j$ and $h \in [H]$, $f_i(S_{h-1}) < 1$ and that for any $i \notin R_j^*$, $f_i(T^*) = 1$. Hence by summing the inequality in the above proposition over all functions f_i in $R_j \setminus R_j^*$, we have

$$\Delta_j \ \geq \ \frac{1}{\rho} \sum_{h=1}^{H} \ \sum_{i \in R_j \setminus R_j^*} \frac{f_i(T^*) - f_i(S_{h-1})}{1 - f_i(S_{h-1})} \ \geq \ \frac{H}{\rho}(|R_j| - |R_j^*|) \ = \ 4\alpha(|R_j| - |R_j^*|) \tag{4}$$

We now upper bound Δ_j. Note by definition that for any $i \notin R_{j-1}$, $f_i(S_0) = 1$ and therefore f_i does not contribute to Δ_j. For any $i \in R_{j-1}$, the total contribution of f_i to Δ_j is at most α by Claim 3.4. Hence it follows that

$$\Delta_j \leq \alpha |R_{j-1}| \tag{5}$$

Combining (4) and (5) completes the proof of Lemma 4.5.

5 New LP Relaxation for Covering Steiner Tree

We revisit the *covering Steiner tree* problem (CST) [20, 23]. Recall that its input consists of a metric (V, d) with root r and a collection of groups $\mathcal{G} \subseteq 2^V$ where each group $g \in \mathcal{G}$ is associated with a requirement k_g. The goal is to find a minimum cost r-rooted tree τ that includes r and covers all groups. A group g is covered if τ contains at least k_g nodes from g.

Konjeved et al. [23] gave an $O((\log m + \log k_{max}) \cdot \log g_{max} \cdot \log n)$-approximation where m is the number of groups, $k_{max} := \max_{g \in \mathcal{G}} k_g$ is the maximum requirement and $g_{max} = \max_{g \in \mathcal{G}} |g|$ is the maximum group size. Later, Gupta and Srinivasan [20] gave an improved $O(\log m \cdot \log g_{max} \cdot \log n)$-approximation that has no dependence on k_{max}. In certain settings of parameters, the result [20] yields an improvement of $\Theta(\frac{\log n}{\log \log n})$ over [23]. Both these previous works use the same LP that has a large integrality gap of $\Omega(k_{max})$ [23]. To circumvent this, they solved a sequence of LPs adaptively: In each step, the requirement of each group is partially satisfied (by a constant fraction), and the subsequent LP is defined over the residual requirement and the unvisited vertices.

We introduce stronger constraints, that yield an LP for CST with integrality gap $O(\log m \cdot \log g_{max} \cdot \log n)$. This new LP is an important ingredient in our algorithm for LCST (next section), and might be useful in other contexts as well.

5.1 Simplifying assumptions

Following [23, 20], WLOG, we will proceed our analysis under the following simplifying assumptions.

- The given graph is a (weighted) tree $T = (V, E)$ with root r and weight w_e on each edge $e \in E$.
- The groups in $\mathcal{G}$ are disjoint.
- Every vertex in a group is a leaf, i.e. has degree 1.
- Every vertex of degree 1 is in a group.

We note that by the tree embedding in [14], a ρ-approximation on tree instances will immediately imply an $O(\rho \log |V|)$-approximation on general instances. Hence we focus on giving an $O(\log m \cdot \log g_{max})$-approximation on tree instances.

5.2 Linear Program

We let L denote the set of leaves in V. Because of the above simplifying assumptions, we can label each vertex v in a group with a unique leaf-edge incident on it, and vice versa. We abuse notation by allowing $j \in L$ to denote both the leaf-vertex and its unique incident edge. For any subset of leaves $L' \subseteq L$, let $\texttt{cut}(r, L')$ denote the family of edge-cuts that separate the root r from all vertices in L'. For notational convenience, we assume that there is a pseudo-edge incident on r that is an ancestor of all edges in E. We call this edge super-root and assume that it is alway included in our solution and has zero cost. Hence for any edge $j \in E$, its parent edge, denoted as $pe(j)$, is well-defined.

We formulate the following linear programming relaxation for CST on tree instances.

$$\min \quad \sum_{e \in E} w_e x_e \tag{$\mathsf{LP}_{\mathsf{CST}}$}$$

$$\text{subject to:} \quad x_{pe(e)} \;\geq\; x_e \qquad \forall\, e \in E \tag{6}$$

$$(k_g - |A|) \sum_{j \in B \setminus L} x_j + \sum_{j \in B \cap (L \setminus A)} x_j \;\geq\; k_g - |A| \qquad \forall g \in \mathcal{G},\ \forall A \subseteq g, \quad \forall B \in \texttt{cut}(r, g \setminus A) \tag{7}$$

$$x_e \;\in\; [0, 1] \qquad \forall e \in E$$

We first argue that this is a valid relaxation. Consider any instance of CST on trees and a fixed feasible solution (tree) τ^*, which gives a natural integral solution: $x_e = 1$ if and only if $e \in \tau^*$. We focus on constraints (7), since the other constraints are obviously satisfied. Consider any $g \in \mathcal{G}$, $A \subseteq g$ and $B \in \texttt{cut}(r, g \setminus A)$. Let $\tau^*(E \setminus A)$ denote the subtree $\tau^* \bigcap (E \setminus A)$. Note that $\tau^*(E \setminus A)$ is connected, since A consists only of leaf edges. Since τ^* has at least k_g edges in g (it is a feasible CST solution), $|\tau^*(E \setminus A) \bigcap (g \setminus A)| \geq k_g - |A|$.

- Suppose that there exists $\overline{j} \in \tau^*(E \setminus A) \cap B$ such that $\overline{j} \notin L$. Then since $\overline{j} \in B \setminus L$, it follows that $(k_g - |A|) \sum_{j \in B \setminus L} x_j \geq k_g - |A|$, hence the constraint is satisfied.
- The remaining case has $\tau^*(E \setminus A) \cap B \subseteq L$. In words, B cuts $g \setminus A$ from r using only leaf edges; so $B \supseteq \tau^*(E \setminus A) \bigcap (g \setminus A)$. Thus $\sum_{j \in B \cap (L \setminus A)} x_j \geq |\tau^*(E \setminus A) \bigcap (g \setminus A)| \geq k_g - |A|$.

In either of the above cases, constraint (7) is satisfied.

5.3 Solving the LP

Since $\mathsf{LP}_{\mathsf{CST}}$ has exponentially many constraints, in order to solve it in polynomial time, we need a separation oracle. Again we will focus on constraints (7), since other constraints are only polynomially many. We observe such an oracle is reduced to finding a polynomial time algorithm for the following problem.

Problem $\mathsf{MinCutWithExceptions}$: In this problem, we are given as input a tree T rooted at r with leaves L and cost $\ell(e)$ on each edge e. Further, an integer $D \geq 0$ is given. The goal is to find a minimum cost cut that separates r from any D leaves.

We first explain how the desired separation oracle is reduced to solving the above problem. Consider the constraints (7) for a fixed $g \in \mathcal{G}$ and all $A \subseteq g$ with $|A| = \eta$ (finally we will iterate over all $g \in \mathcal{G}$ and $0 \leq \eta \leq n$). Then $k_g - |A|$ (the right-hand-side of the constraints) is also fixed. Given x_j values, we would like to find $A \subseteq g$ with $|A| = \eta$ and $B \in \mathtt{cut}(r, g \setminus A)$ that minimizes the left-hand-side, and test if this is smaller than $k_g - \eta$. Formally, we can recast this into $\mathsf{MinCutWithExceptions}$ as follows: Remove all edges from E that are not on any path from the root r to a vertex in g, and let T' be the resulting tree and this is the input tree to the problem. Note that leaves of T' are precisely g. For all leaf-edges $j \in g$, let $\ell(j) := x_j$; and for all non-leaf $j \in T' \setminus g$, $\ell(j) := (k_g - \eta) \cdot x_j$. Also set bound $D := \eta - |g|$.

We next show that $\mathsf{MinCutWithExceptions}$ can be solved via a dynamic programming.

Lemma 5.1. *The problem $\mathsf{MinCutWithExceptions}$ can be solved in polynomial time.*

Proof. To formally describe our dynamic programming, we make some simplifying assumptions. By introducing dummy edges of infinite cost, we assume WLOG that the tree T is binary and the root r is incident to exactly one edge e_r. Hence every non-leaf edge e has exactly two child-edges e_1 and e_2. For any edge $e \in T$ let T_e denote the subtree of T rooted at e, i.e. T_e contains edge e and all its descendants.

We define a recurrence for $C[e, k]$ which denotes the minimum cost cut that separates the root of T_e from exactly k leaves in T_e. Note that $C[e_r, D]$ gives the optimal value.

For any leaf-edge f set:

$$C[f, k] = \begin{cases} 0 & \text{if } k = 0, \\ \ell(f) & \text{if } k = 1, \text{ and} \\ \infty & \text{otherwise.} \end{cases}$$

For any non-leaf edge e with children e_1 and e_2, set:

$$C[e, k] = \begin{cases} 0 & \text{if } k = 0, \\ \min\{C[e_1, k_1] + C[e_2, k_2] \,:\, k_1 + k_2 = k\} & \text{if } 1 \leq k < |L \cap T_e|, \\ \min\{\ell(e),\, \min\{C[e_1, k_1] + C[e_2, k_2] \,:\, k_1 + k_2 = k\}\} & \text{if } k = |L \cap T_e|, \text{ and} \\ \infty & \text{otherwise.} \end{cases}$$

It can be checked directly that this recurrence computes the desired values. □

5.4 Algorithm

Our algorithm uses a threshold rounding and then applies the elegant rounding scheme from [23] (to a suitably modified fractional solution). In this respect it is somewhat similar to the algorithm in [20]; however that approach solves LPs and rounds them iteratively. Whereas our algorithm uses a single LP. We describe the KRS rounding scheme in Section 5.4.1 and our main algorithm in Section 5.4.2.

5.4.1 The rounding scheme in [23]

Algorithm 3 ALG-KRS [23]

```
INPUT: Undirected tree $T = (V, E)$ rooted at $r$; $z_e \in [0,1], \forall e \in E$ such that for all $e \in E$, $z_{pe(e)} \geq z_e$.
1: $S \leftarrow \emptyset$.
2: For each $e \in E$ such that $z_e > 0$, add $e$ to $S$ with probability $\frac{z_e}{z_{pe(e)}}$.
OUTPUT: The connected component (tree) including $r$ in $S$.
```

The following proposition easily follows.

Proposition 5.2 ([23]). *Each edge e is included in the final solution of ALG-KRS with probability z_e.*

Definition 5.3 (KRS properties). *For all $g \in \mathcal{G}$, consider any $R(g) \subseteq g$ and any $0 \leq r_g \leq |R(g)|$. We say that $(z, R(g), r_g)$ satisfies the KRS properties if it satisfies the following properties.*

$$z_{pe(e)} \geq z_e \qquad \forall e \in E \tag{8}$$

$$\sum_{j \in T(e) \cap R(g)} z_j \leq r_g z_e \qquad \forall e \in E \tag{9}$$

where $T(e)$ is the subtree below (and including) edge e.

The first property (8) is the same as the constraints (6). The second property (9) is a Lipschitz-type condition that says conditional on any edge e being chosen, its subtree $T(e)$ can contribute at most r_g to the requirement of $R(g)$. We remark that this property (9) was enforced explicitly in [23, 20] as constraints (for specific $R(g)$ and r_g) in their sequence of LPs. On the other hand, as we show later, our LP constraint (7) directly gives this property for all relevant $R(g)$ and r_g values.

Lemma 5.4 ([23]). *Suppose that $(z, R(g), r_g)$ satisfies the KRS properties. Let L_{krs} denote the set of leaves that are covered by ALG-KRS with the input $z_e, e \in E$. Consider any constant $\delta \in [0,1]$. Then for any $g \in \mathcal{G}$, we have*

$$\Pr\Big[|L_{krs} \cap R(g)| \leq (1-\delta)\mu_g\Big] \leq \exp\Big(-\frac{\delta^2 \cdot \mu_g}{2 + r_g(1 + \ln |R(g)|)}\Big)$$

where $\mu_g := \mathbb{E}\big[|L_{krs} \cap R(g)|\big] = \sum_{j \in R(g)} z_j$.

Proof. We only give a quick sketch of the proof, since this is implicit in [23]. For any $j, j' \in R(g)$, we say that $j \sim j'$ if and only if (1) $j \neq j'$ and (2) the least common ancestor $\texttt{lca}(j,j')$ of j and j' is not r. Define

$$\Delta_g := \sum_{j,j' \in R(g): j \sim j', z_{\texttt{lca}(j,j')} > 0} \frac{z_j \cdot z_{j'}}{z_{\texttt{lca}(j,j')}}$$

In Theorem 3.2 in [23], Konjevod et al. showed using the KRS properties that

$$\Delta_g \leq \mu_g(r_g - 1 + r_g \ln |R(g)|)$$

We note that the proof of Theorem 3.2 implies this, although it is stated only for $\mu_g = r_g$. Further, they used this bound in Jansen's inequality to obtain for any $\delta \in [0,1]$,

$$\Pr\Big[|L_{krs} \cap R(g)| \leq (1-\delta)\mu_g\Big] \leq \exp\Big(-\frac{\delta^2 \mu_g}{2 + \Delta_g/\mu_g}\Big)$$

Above $\mu_g = \mathbb{E}\big[|L_{krs} \cap R(g)|\big] = \sum_{j \in R(g)} \Pr[j \in L_{krs}] = \sum_{j \in R(g)} z_j$, by Proposition 5.2. Combining the above two inequalities yields the lemma. ☐

5.4.2 Main algorithm

The algorithm first preprocesses x^* an optimal solution to $\mathsf{LP}_{\mathsf{CST}}$ to obtain another solution x' that will used in the KRS-rounding. Let $E^* := \{e \in E \mid x^*_e \geq \frac{1}{2}\}$; all edges in E^* will be included in our solution with probability 1. Observe that E^* forms a tree rooted at r due to the constraints (7). We let $T' = T/E^*$ denote the tree obtained by contracting edges E^* in the original tree T; and $E' = E \setminus E^*$ the edges of T'. For any group $g \in \mathcal{G}$, define $R(g) := g \setminus E^*$ and $r_g := k_g - |g \cap E^*|$. Since E^* already covers the vertices in $g \cap E^*$, we will henceforth be concerned with getting coverage of at least r_g from vertices $R(g)$. Our algorithm ALG-CST is formally described below. Note that for any $e \in E'$, $z'_e = 2x'_e \leq 1$ and $z'_{pe(e)} \geq z'_e$, hence z' is a valid input to ALG-KRS with respect to T'. The final solution is the union of all sampled trees τ plus the edges in E^* that are initially shrunk.

Algorithm 4 ALG-CST

```
INPUT: Tree $T$ with edge lengths, root $r$, groups $\mathcal{G}$ and requirements $\{k_g\}_{g\in\mathcal{G}}$.
Let $x^*_e$, $e \in E$ be an optimal solution to $\mathsf{LP}_{\mathsf{CST}}$.
$E^* \leftarrow \{e \in E \mid x^*_e \geq 1/2\}$, $R(g) \leftarrow g \setminus E^*$ and $r_g \leftarrow k_g - |g \cap E^*|$.
Shrink all edges in $E^*$ and let $T'$ be the resulting tree with edges $E' := E \setminus E^*$.
Obtain solution $x'$ from $x^*$ using the procedure in Lemma 5.5.
For each $e \in E'$, $z'_e \leftarrow 2x'_e$; note $z'_e \in [0, 1]$.
$\mathcal{S} \leftarrow \emptyset$.
Repeat the following $\theta(\log m \cdot \log g_{max})$ times:
    Add to $\mathcal{S}$ the tree $\tau$ produced by ALG-KRS with $z_e = z'_e$ for tree $T'$.
OUTPUT: Combine all trees in $\mathcal{S}$ and $E^*$ and output the resulting tree.
```

Lemma 5.5. *There exist $x'_e \in [0, x^*_e]$ for all $e \in E \setminus E^*$ such that for all $g \in \mathcal{G}$, $(x', R(g), r_g)$ satisfies the KRS properties in tree T' and $\sum_{j\in R(g)} x'_j \geq r_g$. Furthermore, such x' can be computed in polynomial time.*

Proof. Consider constraints (7) of $\mathsf{LP}_{\mathsf{CST}}$. Fix a group $g \in \mathcal{G}$ and let $A := g \cap E^*$. Consider tree T' as a flow network with each leaf edge f having capacity x^*_f and each non-leaf edge e having capacity $r_g \cdot x^*_e$. The root r is the source and leaves $R(g) = g \setminus A$ are the sinks. Then constraints (7) imply that the min cut separating r from $R(g)$ has value at least r_g: note that although these constraints are for the original tree T, they imply similar constraints for T' since T' is obtained from T by edge-contraction.[4] Hence there must exist a max-flow of volume at least r_g from r to $R(g)$ in the above network. Let x'_f denote the volume of this flow into each leaf edge $f \in R(g)$; clearly we have that $x'_f \leq x^*_f$ (due to capacity on leaves) and:

$$\sum_{j\in R(g)} x'_j \geq r_g. \tag{10}$$

Moreover, by the capacities on non-leaves,

$$\sum_{j\in T(e)\cap R(g)} x'_j \leq r_g x^*_e, \quad \forall e \in E \tag{11}$$

---

[4] In particular every cut B' separating r from $g \setminus A$ in T' is also a cut separating r from $g \setminus A$ in T.

We can use the above procedure on each group $g \in \mathcal{G}$ separately to compute x'_f for all leaf edges f; this is well-defined since groups are disjoint. For each non-leaf edge $e \in E'$ set $x'_e := x^*_e$. Thus we have $0 \leq x'_e \leq x^*_e$ for all $e \in E'$. Observe that this computation can easily be done in poly-time.

Now, (11) implies the second KRS property (9). Property (8) follows, since for each $e \in E'$, we have $x'_{pe(e)} = x^*_{pe(e)} \geq x^*_e \geq x'_e$; the first inequality is due to constraint (6) of $\mathsf{LP}_{\mathsf{CST}}$. Finally (10) implies the coverage property claimed in the lemma. □

5.5 Analysis

Consider a random tree τ in Line (8) of ALG-CST. For a formal analysis, we assume that ALG-CST samples τ, $4\log m(3 + \ln g_{max})$ times independently. For each group $g \in \mathcal{G}$, we will focus on lower bounding the probability that τ satisfies the residual requirement r_g for $R(g)$.

Lemma 5.6. *For any group $g \in \mathcal{G}$, $\Pr[|\tau \cap R(g)| < r_g] \leq \exp(-\frac{1}{2(3+\ln g_{max})})$.*

Proof. Consider any group $g \in \mathcal{G}$. By Lemma 5.5, we know $(x', R(g), r_g)$ satisfies the KRS properties with respect to tree T'. From the definition of $z'_e = 2 \cdot x'_e$ and that $z'_e \leq 1$ for all $e \in E'$, it immediately follows that $(z', R(g), r_g)$ also satisfies the KRS properties on T'. Further, by Proposition 5.2, it follows

$$\mu_g := \mathbb{E}[|\tau \cap R(g)|] = \sum_{j \in R(g)} z'_j = 2 \sum_{j \in R(g)} x'_j \geq 2r_g,$$

where we used the coverage property $\sum_{j \in R(g)} x'_j \geq r_g$ given by Lemma 5.5. By applying Lemma 5.4 with $\delta = 1/2$, we have

$$\begin{aligned} & \Pr\Big[|\tau \cap R(g)| < r_g\Big] \leq \exp\Big(-\frac{r_g}{2(2 + r_g(1 + \ln|R(g)|))}\Big) \\ = \; & \exp\Big(-\frac{1}{2(3 + \ln|R(g)|}\Big) \leq \exp\Big(-\frac{1}{2(3 + \ln g_{max})}\Big) \end{aligned}$$

Above we used $r_g \geq 1$ and $|R(g)| \leq g_{max}$. □

It follows that upon sampling the tree τ independently $4 \ln m \cdot (3 + \ln g_{max})$ times, each group $g \in \mathcal{G}$ is covered with probability at least $1 - \frac{1}{m^2}$. A simple union bound yields that all groups are covered by the output of ALG-CST with probability at least $1 - 1/m$.

Now we upper bound the expected cost of the final solution.

Lemma 5.7. *With at least a half probability, the final solution of ALG-CST does not cost more than $8 \log m(3 + \ln g_{max})$ times the optimum.*

Proof. From Proposition 5.2, we know that the expected cost of τ is $\sum_{e \in E \setminus E^*} z'_e \leq 2\sum_{e \in E \setminus E^*} x'_e \leq 2\sum_{e \in E \setminus E^*} x^*_e$. Hence, knowing that $x^*_e \geq 1/2$ for all $e \in E^*$, the expected cost of the final solution is at most $4\log m(3 + \ln g_{max}) \sum_{e \in E \setminus E^*} w_e x^*_e + \sum_{e \in E^*} w_e \leq 4 \log m(3 + \ln g_{max})\mathsf{OPT}$. Markov inequality yields the lemma. □

Hence we derive the following theorem.

Theorem 5.8 ([20]). *There is a randomized $O(\log m \cdot \log g_{max})$-approximation algorithm for CST on tree metrics. For general metrics, there is a randomized $O(\log m \cdot \log g_{max} \cdot \log|V|)$-approximation algorithm.*

6 Latency Covering Steiner Tree Problem

In this section, we give an improved algorithm for the *latency covering Steiner tree problem* (LCST), which is an interesting special case of MLSC. Recall that the input to LCST consists of a symmetric metric (V, d), root $r \in V$ and a collection $\mathcal{G}$ of m groups with associated requirements $\{k_g\}_{g \in \mathcal{G}}$.The goal is find a path staring from r that minimizes the total cover time of all groups. We say that group g is covered at the earliest time t when the path within distance t visits at least k_g vertices in g. We give an $O(\log g_{max} \cdot \log |V|)$-approximation algorithm for this problem where $g_{max} := \max_{g \in \mathcal{G}} |g|$ is the maximum group size.

The stronger constraints (7), which we introduced in the previous section, will play a crucial role in our algorithm and analysis for LCST as well. We make the same simplifying assumptions that we made in Section 5.1. In particular, in order to obtain an $O(\log g_{max} \cdot \log |V|)$-approximation on general instances, it suffices to give an $O(\log g_{\max})$-approximation on tree instances. We will follow most of the notation and will explain any difference in notation if and when necessary.

6.1 Linear Programming Relaxation

We formulate the following linear relaxation for tree instances of LCST.

$$\min \quad \frac{1}{2} \cdot \sum_{\ell \geq 0} 2^\ell \sum_{g \in \mathcal{G}} (1 - y_g^\ell) \tag{$\mathsf{LP}_{\mathsf{LCST}}$}$$

$$\text{subject to:} \quad x_{pe(e)}^\ell \geq x_e^\ell \qquad \forall\, \ell \geq 0, e \in E \tag{12}$$

$$\sum_{j \in E} w_e x_e^\ell \leq 2^\ell \qquad \forall\, \ell \geq 0 \tag{13}$$

$$(k_g - |A|) \sum_{j \in B \setminus L} x_j^\ell + \sum_{j \in B \cap L \setminus A} x_j^\ell \geq (k_g - |A|) \cdot y_g^\ell$$

$$\forall \ell \geq 0, g \in \mathcal{G}, A \subseteq g,\ B \in \mathtt{cut}(r, g \setminus A) \tag{14}$$

$$y_g^{\ell+1} \geq y_g^\ell \qquad \forall\, \ell \geq 0, g \in \mathcal{G} \tag{15}$$

$$x_e^\ell \in [0,1] \qquad \forall\, \ell \geq 0, e \in E$$

$$y_g^\ell \in [0,1] \qquad \forall\, \ell \geq 0, g \in \mathcal{G}$$

To see that this is a valid relaxation, let OPT denote the optimal path. For any $\ell \geq 0$ let $\mathsf{OPT}(2^\ell)$ denote the prefix of length 2^ℓ of OPT. We construct a feasible integral solution for OPTas follows. The variable x_e^ℓ indicates if edge e is covered by $\mathsf{OPT}(2^\ell)$. The indicator variable y_g^ℓ has value one if and only if group g is covered by $\mathsf{OPT}(2^\ell)$. Constraints (12) come from the fact that $\mathsf{OPT}(2^\ell)$ is a path starting at r. Constraints (13) say that the edges in $\mathsf{OPT}(2^\ell)$ have a total weight of at most 2^ℓ, which is clearly true. Note that for each $\ell \geq 0$, we created a set of constraints (14) that are essentially equivalent to the constraints (7) of $\mathsf{LP}_{\mathsf{CST}}$. Constraints (14) can be shown to be valid similarly as are the constraints (7). Constraints (15) enforce the fact that if group g is covered by $\mathsf{OPT}(2^\ell)$ and it must be covered by $\mathsf{OPT}(2^{\ell+1})$ as well. Now consider the objective value: the total contribution of a group g that is covered by OPTat time $\in (2^k,\, 2^{k+1}]$ is $\frac{1}{2} \cdot \sum_{\ell=0}^{k} 2^\ell \leq 2^k$. Thus the objective value of this integral solution is at most OPT.

It is easy to see that it suffices to consider ℓ of polynomial value: so the number of variables is polynomial. Note that constraints (14) are exponentially many. However, for each fixed ℓ, we can use essentially the same separation oracle that we used for the constraints (7) of $\mathsf{LP}_{\mathsf{CST}}$.

6.2 Algorithm

Let $(\overline{x}, \overline{y})$ denote a fixed optimal solution to $\mathsf{LP}_{\mathsf{LCST}}$. Let $\overline{E}^{\ell} := \{e \in E \mid \overline{x}_e^{\ell} \geq 1/4\}$. Observe that $\overline{E}^{\ell}$ forms a tree rooted at r due to the constraints (12). The edges in $\overline{E}^{\ell}$ will be added to our solution with probability 1 by the ℓth phase of our algorithm. Let $R^{\ell}(g) := g \setminus \overline{E}^{\ell}$ and $r_g^{\ell} = k_g - |g \cap \overline{E}^{\ell}|$. Our algorithm will try to satisfy the residual requirement r_g^{ℓ} of group g by covering vertices in $R^{\ell}(g)$ when the group is substantially covered (say to extent $\frac{1}{2}$) by the optimal fractional solution $(\overline{x}, \overline{y})$.

For each $\ell \geq 0$ we preprocess $\overline{x}^{\ell}$ and obtain $\tilde{x}^{\ell}$ that have "good" properties.

Lemma 6.1. *For any $\ell \geq 0$, we can find in polynomial time $\tilde{x}_e^{\ell} \in [0, \overline{x}_e^{\ell}]$, $\forall e \in E \setminus \overline{E}^{\ell}$ such that $\forall g \in \mathcal{G}$:*

1. $(\tilde{x}, R^{\ell}(g), r_g^{\ell})$ *satisfies the KRS-properties in tree $\tilde{T}^{\ell}$.*
2. $\sum_{j \in R^{\ell}(g)} \tilde{x}_j^{\ell} \;\geq\; r_g^{\ell} \cdot \overline{y}_g^{\ell}$ *(coverage property).*

Proof. The proof of this lemma is very similar to Lemma 5.5. For each ℓ, we obtain the desired $\tilde{x}^{\ell}$ values from $\overline{x}^{\ell}$ by taking similar steps as we did to set x' from x^* in the proof of Lemma 5.5. ☐

We now describe our algorithm ALG-LCST as follows. We will say that ALG-LCST is in phase ℓ based on the value of ℓ in the description.

Algorithm 5 ALG-LCST

```
INPUT: Tree $T$ with edge lengths, root $r$, groups $\mathcal{G}$ and requirements $\{k_g\}_{g \in \mathcal{G}}$.
$\pi \leftarrow \emptyset$.
Let $(\overline{x}, \overline{y})$ be an optimal solution to $\mathsf{LP}_{\mathsf{LCST}}$.
for $\ell = 0, 1, 2, ...$ do
    $\overline{E}^{\ell} \leftarrow \{e \in E \mid \overline{x}_e \geq 1/4\}$, $R^{\ell}(g) \leftarrow g \setminus \overline{E}^{\ell}$ and $r_g \leftarrow k_g - |g \cap \overline{E}^{\ell}|$.
    Shrink all edges in $\overline{E}^{\ell}$ in $T$ and let $\tilde{T}^{\ell}$ be the resulting tree with the edge set $\tilde{E}^{\ell} := E \setminus \overline{E}^{\ell}$.
    Obtain solution $\tilde{x}^{\ell}$ from $\overline{x}^{\ell}$ using Lemma 6.1.
    For each $e \in \tilde{E}^{\ell}$, $z_e^{\ell} \leftarrow 4\tilde{x}_e^{\ell}$; note that $z_e^{\ell} \in [0, 1]$.
    $\mathcal{S}^{\ell} \leftarrow \emptyset$.
    Repeat the following $6(3 + \log g_{max})$ times:
        $\tau^{\ell} \leftarrow$ the tree produced by ALG-KRS with $z = z^{\ell}$ on tree $\tilde{T}^{\ell}$
        Add $\tau^{\ell}$ to $\mathcal{S}^{\ell}$
    Combine all trees in $\mathcal{S}^{\ell}$ with $\overline{E}^{\ell}$ and make the resulting tree as a tour $P^{\ell}$ starting at $r$
    If $P^{\ell}$ has weight at most $192(3 + \log g_{max}) \cdot 2^{\ell}$ then
        $\pi \leftarrow \pi \cdot P^{\ell}$.
end for
OUTPUT: The path $\pi$.
```

6.3 Analysis

For any group g, define $\ell(g)$ as the smallest $\ell \geq 0$ such that $\overline{y}_g^{\ell} \geq 1/2$. Then it follows that for any $\ell \geq \ell(g)$, $\overline{y}_g^{\ell} \geq 1/2$ due to constraints (15) of $\mathsf{LP}_{\mathsf{LCST}}$. In words, the optimal fractional solution covers group g by at least half within time $2^{\ell(g)}$. Consider any group $g \in \mathcal{G}$, $\ell \geq \ell(g)$ and a tree τ^ℓ in Line (10) of ALG-LCST. Since all edges in $\overline{E}^\ell$ are included in P^ℓ with probability 1, group g is covered by P^ℓ if and only if at least r_g^ℓ vertices in the residual group $R^\ell(g)$ are covered by τ^ℓ. This motivates us to derive the following lemma.

Lemma 6.2. *For any $g \in \mathcal{G}$ and $\ell \geq \ell(g)$, $\Pr[|\tau^\ell \cap R(g)| < r_g] \leq \exp(-\frac{1}{2(3+\ln g_{max})})$.*

Proof. From Lemma 6.1 it follows that $(\tilde{x}^\ell, R^\ell(g), r_g^\ell)$ satisfies the KRS properties on tree $\tilde{T}^\ell$. Furthermore, using $\overline{y}_g^\ell \geq \frac{1}{2}$ and the coverage property in Lemma 6.1:

$$\mu_g^\ell \;:=\; \mathbb{E}[|\tau^\ell \cap R^\ell(g)|] \;=\; \sum_{j \in R^\ell(g)} z_j^\ell \;=\; 4 \cdot \sum_{j \in R^\ell(g)} \tilde{x}_j^\ell \;\geq\; 4 \cdot r_g^\ell \cdot \overline{y}_g^\ell \;\geq\; 2 r_g^\ell$$

Here we also used Proposition 5.2 that $\Pr[j \in \tau^\ell] = z_j^\ell$. By applying Lemma 5.4 with $\delta = 1/2$, we have

$$\Pr\Big[|\tau^\ell \cap R(g)| < r_g\Big] \leq \exp\Big(-\frac{r_g^\ell}{2(2 + r_g^\ell(1 + \ln|R^\ell(g)|))}\Big) \leq \exp\Big(-\frac{1}{2(3 + \ln g_{max})}\Big)$$

This proves Lemma 6.2. ☐

We now show the following lemma.

Lemma 6.3. *Consider any group $g \in \mathcal{G}$ and $\ell \geq \ell(g)$. The probability that P^ℓ has a total weight of at most $192(3 + \log g_{max}) \cdot 2^\ell$ and covers g is at least $3/4$.*

Proof. By Proposition 5.2, we know that each edge $e \in \tilde{E}^\ell$ is included in τ^ℓ with probability $z_e^\ell = 4x'^\ell_e$. Since for all $e \in \overline{E}^\ell$, $\overline{x}_e \geq 1/4$, the expected total weight of the edges in $\overline{E}^\ell$ and τ^ℓ is upper bounded by

$$\sum_{e \in \overline{E}^\ell} w_e + \sum_{e \in \tilde{E}^\ell} w_e \cdot 4\tilde{x}^\ell \quad \leq \quad 4 \sum_{e \in E} w_e \cdot \overline{x}_e^\ell \quad \leq \quad 4 \cdot 2^\ell$$

The last inequality is due to the constraints (13). Hence the expected cost of P^ℓ is at most $24(3 + \log g_{max}) \cdot 2^\ell$. Markov inequality immediately gives that the total weight of P^ℓ is greater than $192(3 + \log g_{max}) \cdot 2^\ell$ with probability at most $1/8$. Since τ_ℓ is sampled $6(3 + \log g_{max}) \cdot 2^\ell$ times, from Lemma 6.2, we know that group g is not covered by P^ℓ with a probability of at most $1/e^3 \leq 1/8$. Hence the claim follows. ☐

Fix any group $g \in \mathcal{G}$, and $\ell \geq \ell(g)$. Among $P^{\ell(g)}$, $P^{\ell(g)+1}$, ... , P^ℓ, consider the paths that are added to π. Clearly the total weight of such paths is at most $O(\log g_{max} \cdot 2^\ell)$. By Lemma 6.3, the probability that none of these paths covers g is at most $\frac{1}{4^{\ell - \ell(g)+1}}$. Hence the expected cover time of g is at most $\sum_{\ell \geq \ell(g)} O(\log g_{max}) \cdot 2^\ell \cdot \frac{1}{4^{\ell-\ell(g)+1}} = O(\log g_{max} \cdot 2^{\ell(g)})$. Thus the expected total cover time

$$\textsf{ALG-LCST} \quad \leq \quad O(\log g_{max}) \cdot \sum_{g \in \mathcal{G}} 2^{\ell(g)}$$

By the definition of $2^{\ell(g)}$ being $\frac{1}{2}$-completion time in the LP, we know

$$\textsf{OPT} \;\geq\; \frac{1}{2} \cdot \sum_{\ell \geq 0} 2^\ell \sum_{g \in \mathcal{G}} (1 - \overline{y}_g^\ell) \;\geq\; \frac{1}{2} \sum_{g \in \mathcal{G}} 2^{\ell(g)-1} \Big(1 - \overline{y}_g^{\ell(g)-1}\Big) \;\geq\; \frac{1}{8} \cdot \sum_{g \in \mathcal{G}} 2^{\ell(g)}$$

These two inequalities show that ALG-LCST is an $O(\log g_{max})$-approximation on tree instances. By the tree embedding due to [14], we conclude that ALG-LCST can be used to yield an $O(\log g_{max} \cdot \log|V|)$-approximation for general metrics, thus proving Theorem 1.3.

7 Stochastic Submodular Ranking

In this section, we study the *stochastic submodular ranking* problem. The input consists of a set $\mathcal{A} = \{X_1, ..., X_n\}$ of n independent random variables (stochastic elements, each over domain Δ) with integer sizes $\{\ell_i\}_{i=1}^n$ (deterministic), and m monotone submodular functions $f_1, ..., f_m : 2^\Delta \to [0, 1]$ on groundset Δ. We are also given the distribution (over Δ) of each stochastic element $\{X_i\}_{i=1}^n$. The realization $x_i \in \Delta$ of random variable X_i is known immediately after scheduling/displaying X_i. Here, X_i requires ℓ_i units of time to be scheduled; recall that in the non-stochastic submodular ranking, elements were assumed to have unit sizes. A feasible solution/policy is an adaptive ordering of $\mathcal{A}$, represented naturally by a decision tree with branches corresponding to the realization of the stochastic elements. We use $\langle \pi(1), \ldots, \pi(n) \rangle$ to denote this ordering, where each $\pi(l)$ is a random variable denoting the index of the l^{th} scheduled element.

The cover time $\mathtt{cov}(f_i)$ of f_i is defined as the earliest time t such that f_i has value one on the realization of the elements completely scheduled within time t. More formally, $\mathtt{cov}(f_i)$ is the earliest time t such that $f_i(\{x_{\pi(1)}, ..., x_{\pi(k_t)}\}) = 1$ where k_t is the maximum integer such that $\ell_{\pi(1)} + \ell_{\pi(2)} + ... + \ell_{\pi(k_t)} \leq t$. If the function value never reaches one (due to the stochastic nature of elements) then $\mathtt{cov}(f_i) = \ell_1 + \ell_2 + ... + \ell_n$. Note that the cover time is a random value. The goal is to find a policy that (approximately) minimizes the expected total cover time $\mathbb{E}\left[\sum_{i \in [m]} \mathtt{cov}(f_i)\right]$. As in the previous sections, each function f_i satisfies the following property: for any $S \supseteq S'$, if $f_i(S) - f_i(S') > 0$ then it must be the case that $f_i(S) - f_i(S') \geq \epsilon$, where $\epsilon > 0$ is a uniform value for all functions f_i.

Our stochastic extension captures many interesting applications.

Adaptive Set Cover: We are given as input a ground set Δ, and a collection $\mathcal{S} \subseteq 2^\Delta$ of (non-stochastic) sets. There are stochastic elements $\{X_i : i \in [n]\}$, each defined over Δ, and having costs. The goal is to give an adaptive policy that hits all sets in $\mathcal{S}$ using the minimum cost of (stochastic) elements. This problem was studied in [17, 26, 24]. The problem can be shown as an instance of stochastic SR: There is a single monotone submodular function $f_1(A) := \frac{1}{|\mathcal{S}|} \sum_{S \in \mathcal{S}} \min\{1, |A \cap S|\}$ and $\epsilon = 1/|\mathcal{S}|$.

Shared Filter Evaluation: This problem was introduced by [26], and the result was improved to an essentially optimal solution in [24]. In this problem, there is a collection of independent filters $X_1, X_2, ..., X_n$, each gets evaluated either to `True` or `False`. For each filter $i \in [n]$, we are given the "selectivity" $p_i = \Pr[X_i \text{ is true}]$ and the cost ℓ_i of running the filter. We are also given a collection $\mathcal{Q}$ of queries, where each query Q_j is a conjunction of a subset of queries. We would like to determine each query in $\mathcal{Q}$ to be `True` or `False` by (adaptively) testing filters of the minimum cost. In order to cast this problem as stochastic SR, we use $\Delta = \bigcup_{i=1}^n \{t_i,\ f_i\}$; for each $i \in [n]$, $X_i = t_i$ w.p. p_i, and $X_i = f_i$ with the remaining probability $1 - p_i$. We create one monotone submodular function:

$$f_1(A) := \frac{1}{|\mathcal{Q}|} \sum_{Q_j \in \mathcal{Q}} \min\left\{1,\ |A \cap \{f_i : i \in Q_j\}| \,+\, \frac{1}{|Q_j|} \cdot |A \cap \{t_i : i \in Q_j\}|\right\}$$

We note that the *Shared Filter Evaluation* problem can be studied for a latency type of objective also. In this case, for each query $Q_j \in \mathcal{Q}$, we create a separate submodular function:

$$f_j(A) := \min\left\{1,\ |A \cap \{f_i : i \in Q_j\}| \,+\, \frac{1}{|Q_j|} \cdot |A \cap \{t_i : i \in Q_j\}|\right\}$$

In this case, the stochastic SR problem corresponds precisely to filter evaluation that minimizes the *average time* to answer queries in $\mathcal{Q}$.

Stochastic generalized min-sum set cover: We are given as input a ground set Δ, and a collection $\mathcal{S} \subseteq 2^{\Delta}$ of (non-stochastic) sets with requirement $k(S)$ for each $S \in \mathcal{S}$. There are stochastic elements $\{X_i : i \in [n]\}$, each defined over Δ. Set $S \in \mathcal{S}$ is said to be completed when at least $k(S)$ elements from S have been displayed. The goal is to find an adaptive ordering of $[n]$ so as to minimize the expected total completion time. This can be reduced to stochastic SR by defining function $f^S(A) := \min\{1, \, |A \cap S|/k(S)\}$ for each $S \in \mathcal{S}$; here $\epsilon = 1/k_{max}$ where k_{max} denotes the maximum requirement.

We consider adaptive policies: this chooses at each time $\ell_{\pi(1)} + \ell_{\pi(2)} + ... + \ell_{\pi(k-1)}$, the element $X_{\pi(k)} \in \mathcal{A} \setminus \{X_{\pi(1)}, X_{\pi(2)}, X_{\pi(3)}, ..., X_{\pi(k-1)}\}$ after observing the realizations $x_{\pi(1)}, ..., x_{\pi(k-1)}$. So it can be described as a decision tree. Our main result is an $O(\log \frac{1}{\epsilon})$-approximate adaptive policy. This result is again inspired by our simpler analysis of [1].

Remark 7.1. *Obviously, the stochastic submodular ranking problem generalizes its deterministic counterpart where elements have non-uniform sizes. We note that [1] studied the deterministic submodular ranking only when all elements have unit sizes.*

Remark 7.2. *For stochastic generalized min-sum set cover (defined above), our result implies an $O(\log k_{max})$-approximation to adaptive policies. However for non-adaptive policies, one can obtain a better $O(1)$-approximation algorithm by combining the Sample Average Approximation (SAA) method [22, 9] with $O(1)$-approximations known for the non-stochastic version [3, 25].*

7.1 Algorithm and Analysis

To formally describe our algorithm, we quickly define the probability spaces we will be concerned with. We use $\Omega = \Delta^n$ to denote the outcome space of $\mathcal{A}$. We use the same notation Ω to denote the probability space induced by this outcome space. For any $S \subseteq \mathcal{A}$ and its realization s, let $\Omega(s)$ denote the outcome subspace that conforms to s. We can naturally define the probability space defined by $\Omega(s)$ as follows: The probability that $w \in \Omega(s)$ occurs is $\Pr_{\Omega}[w]/\Pr_{\Omega}[\Omega(s)]$. We also use $\Omega(s)$ to denote this probability space.

The main algorithm is a natural extension of the algorithm by Azar and Gamzu [1], which is given below. In the algorithm, Let $\alpha := 1 + \ln(\frac{1}{\epsilon})$. In the output, $\pi(l)$ denotes the lth element (random variable) in $\mathcal{A}$ that appears in the output.

Algorithm 6 ALG-AG-STO

```
INPUT: $\mathcal{A} = \{X_1, ..., X_n\}$ with $\{\ell_1, ..., \ell_n\}$; $f_i : 2^{\Delta} \to [0,1]$, $i \in [m]$.
S ← ∅.   ($S$ are the elements completly scheduled so far, and $s$ their instantiation.)
while There exists $f_i$ such that $f_i(s) < 1$ do
    Choose element $X_e$ as follows:
        $$X_e = \arg \max_{X_e \in \mathcal{A} \setminus S} \frac{1}{\ell_e} \cdot \mathbb{E}_{\Omega(s)} \left[ \sum_{i \in [m], f_i(s) < 1} \frac{f_i(s \cup \{X_e\}) - f_i(s)}{1 - f_i(s)} \right]$$
    $S \leftarrow S \bigcup \{X_e\}$.
    $\pi(|S|) \leftarrow X_e$. Schedule $X_e$ and observe its realization.
end while
OUTPUT: An adaptive ordering $\pi$ of $\mathcal{A}$.
```

Observe that taking expectation over $\Omega(s)$ in Step 3 is the same as expectation over the distribution of X_e since $X_e \notin S$ and the elements are independent. Also note that this algorithm implicitly defines a decision tree.

Theorem 7.3. *ALG-AG-STO is an $O(\ln(\frac{1}{\epsilon}))$-approximation algorithm for stochastic submodular ranking.*

To simplify notation, without loss of generality, we assume that α is an integer. Let $R(t)$ denote the (random) set of functions that are not satisfied by ALG-AG-STO before time t. Note that the set $R(t)$ includes the functions that are satisfied exactly at time t. Analogously, the set $R^*(t)$ is defined for the optimal decision tree. For notational convenience, we will use $i \in R(t)$ interchangeably with $f_i \in R(t)$. Let $C(t) := \{f_1, ..., f_m\} \setminus R(t)$ and $C^*(t) := \{f_1, ..., f_m\} \setminus R^*(t)$. Note that all the sets $C(\cdot)$, $C^*(\cdot)$, $R(\cdot)$, $R^*(\cdot)$ are stochastic. We have that $\mathsf{ALG} = \sum_{t\in[n]} |R(t)|$ and $\mathsf{OPT} := \sum_{t\in[n]} |R^*(t)|$ and hence ALG and OPT are stochastic quantities. We will show that $\mathbb{E}\mathsf{ALG} = O(\alpha) \cdot \mathbb{E}\mathsf{OPT}$ which suffices to prove Theorem 7.3.

We will be interested in the number of unsatisfied functions at times $\{8\alpha 2^j \; : \; j \in \mathbb{Z}_+\}$ by ALG-AG-STO and the number of unsatisfied functions at times $\{2^j \; : \; j \in \mathbb{Z}_+\}$ by the optimal solution. Let $R_j := R(8\alpha 2^j)$ and $R^*_j = R^*(2^j)$. It is important to note that R_j and R^*_j are concerned with different times, and they are stochastic. For notational simplicity, we let $R_{-1} := \emptyset$.

We will show the following key lemma. Roughly speaking, it says that the expected number of unsatisfied functions by ALG-AG diminishes quickly, unless it is comparable to the expected number of unsatisfied functions in OPT. Once we get this lemma, we can prove Theorem 7.3 similar to the proof of Theorem 3.1.

Lemma 7.4. *For any $j \geq 0$, we have $\mathbb{E}[|R_j|] \leq \frac{1}{4}\mathbb{E}[|R_{j-1}|] + \mathbb{E}[|R^*_j|]$.*

Proof. The lemma trivially holds for $j = 0$, so we consider any $j \geq 1$. For any $t \geq 1$, we use s_{t-1} to denote the set of elements *completely* scheduled by ALG-AG-STO by time $t-1$ along with their instantiations; clearly this is a random variable. Also, for $t \geq 1$ let $\sigma(t) \in [n]$ denote the (random) index of the element being scheduled during time slot $(t-1, t]$.[5] Note that s_{t-1} determines $\sigma(t)$ precisely, but not the instantiation of $X_{\sigma(t)}$. Let $E^*_j \subseteq \mathcal{A}$ be the (stochastic) set of elements that is completely scheduled by the optimal policy within time 2^j. For a certain stochastic set (or elements) S, we will denote its realization under an outcome w as $S(w)$. For example, $X_i(w) \in \Delta$ is the realization of element X_i for outcome w; and $E^*_j(w)$ is the set of first 2^j elements completely scheduled by OPT (under w) along with their realizations.

For any time t and corresponding outcome s_{t-1}, define a set function:

$$f^{s_{t-1}}(D) \; := \sum_{i\in[m], f_i(s_{t-1})<1} \frac{f_i(s_{t-1} \cup D) - f_i(s_{t-1})}{1 - f_i(s_{t-1})}, \qquad \forall D \subseteq \Delta.$$

We will also use $f_i^{s_{t-1}}(D)$ to denote the term inside the above summation.

It is easy to see that the function $f^{s_{t-1}} : 2^{\Delta} \to \mathbb{R}_+$ is submodular. We also define

$$F^{s_{t-1}}(X_e) \; := \; \mathbb{E}_{\, w \leftarrow \Omega(s_{t-1})} \left[f^{s_{t-1}}(X_e(w)) \right], \qquad \forall X_e \in \mathcal{A}. \tag{16}$$

Observe that this is zero for elements $X_e \in s_{t-1}$.

Proposition 7.5. *Consider any time $t \in [n]$ and outcome s_{t-1}. Note that s_{t-1} determines $\sigma(t)$. Then:*

$$\frac{1}{\ell_{\sigma(t)}} \cdot F^{s_{t-1}}(X_{\sigma(t)}) \; \geq \; \frac{1}{\ell_i} \cdot F^{s_{t-1}}(X_i), \qquad \forall X_i \in \mathcal{A}$$

---

[5] Since elements have different sizes, note that $\sigma(t)$ is different from $\pi(t)$ which is the tth element scheduled by ALG-AG-STO.

Proof. At some time $t' \leq t$ (right after s_{t-1} is observed) ALG-AG-STO chose to schedule element $X_{\sigma(t)}$ over all elements $X_i \in \mathcal{A} \setminus s_{t-1}$. By the greedy rule we know that the claimed inequality holds for any $X_i \in \mathcal{A} \setminus s_{t-1}$. Furthermore, the inequality holds for any element $X_i \in s_{t-1}$, since $F^{s_{t-1}}(X_i) = 0$ for such an element. ☐

We now define the *expected gain* by ALG-AG-STO in step t as:

$$G_t \quad := \quad \mathbb{E}_{s_{t-1}}\left[\frac{1}{\ell_{\sigma(t)}} F^{s_{t-1}}(X_{\sigma(t)})\right]. \tag{17}$$

And the expected total gain:

$$\Delta_j \quad := \quad \sum_{t=8\alpha 2^{j-1}}^{8\alpha 2^j} G_t \tag{18}$$

We complete the proof of Lemma 7.4 by upper and lower bounding Δ_j.

Upper bound for Δ_j. Fix any outcome $w \in \Omega$. Below, all variables are *conditioned on* w and hence they are all deterministic. (For ease of notation we do not write w in front of the variables).

$$\begin{aligned}
\Delta_j \quad &:= \quad \sum_{t=8\alpha 2^{j-1}}^{8\alpha 2^j} \frac{1}{\ell_{\sigma(t)}}\, f^{s_{t-1}}(x_{\sigma(t)}) \quad = \quad \sum_{t=8\alpha 2^{j-1}}^{8\alpha 2^j} \frac{1}{\ell_{\sigma(t)}} \quad \sum_{i\in[m]: f_i(s_{t-1})<1} f_i^{s_{t-1}}(x_{\sigma(t)}) \\
&\leq \quad \sum_{t=8\alpha 2^{j-1}}^{8\alpha 2^j} \frac{1}{\ell_{\sigma(t)}} \quad \sum_{i\in R_{j-1}} f_i^{s_{t-1}}(x_{\sigma(t)}) \quad \leq \quad \sum_{t\geq 1} \frac{1}{\ell_{\sigma(t)}} \quad \sum_{i\in R_{j-1}} f_i^{s_{t-1}}(x_{\sigma(t)}) \\
&= \quad \sum_{i\in R_{j-1}} \sum_{k=1}^{n} \frac{f_i(T_k) - f_i(T_{k-1})}{1 - f_i(T_{k-1})}
\end{aligned}$$

The first inequality uses the fact that any $i \notin R_{j-1}$ has f_i already covered before time $8\alpha\, 2^{j-1}$, and so it never contributes to Δ_j. In the last expression, $T_k := \{x_{\pi(1)}, ..., x_{\pi(k)}\} \subseteq \Delta$, the first k instantiations seen under w. The equality uses the fact that for each $\sum_{j=1}^{k-1} \ell_{\pi(j)} < t \leq \sum_{j=1}^{k} \ell_{\pi(j)}$ we have $s_{t-1} = T_{k-1}$ and $\sigma(t) = k$. Finally, by Claim 3.4, the contribution of each function $f_i \in R_{j-1}$ is at most $\alpha := 1 + \ln\frac{1}{\epsilon}$. Thus we obtain $\Delta_j(w) \leq \alpha |R_{j-1}(w)|$, and taking expectations,

$$\Delta_j \quad \leq \quad \alpha \mathbb{E}[|R_{j-1}|] \tag{19}$$

Lower bound for Δ_j. Consider any $8\alpha 2^{j-1} \leq t \leq 8\alpha 2^j$. We will lower bound G_t. Condition on s_{t-1}; this determines $\sigma(t)$ (but not $x_{\sigma(t)}$). Note that $\sum_{i=1}^{n} \ell_i \cdot \Pr[X_i \in E_j^* | s_{t-1}] \leq 2^j$ by definition of E_j^* being the elements that are completely scheduled by time 2^j in OPT. Hence we have

$$\sum_{X_i\in\mathcal{A}} \frac{\ell_i}{2^j} \cdot \Pr[X_i \in E_j^* | s_{t-1}] \leq 1.$$

By applying Proposition 7.5 with the convex multipliers (over i) given above,

$$\begin{aligned}
\frac{1}{\ell_{\sigma(t)}} F^{s_{t-1}}(X_{\sigma(t)}) &\geq \sum_{X_i \in \mathcal{A}} \frac{\ell_i}{2^j} \Pr[X_i \in E_j^* | s_{t-1}] \cdot \frac{1}{\ell_i} F^{s_{t-1}}(X_i) \\
&= \frac{1}{2^j} \sum_{X_i \in \mathcal{A}} \Pr[X_i \in E_j^* | s_{t-1}] \sum_{x_i \in \Delta} \Pr[X_i = x_i | s_{t-1}] \cdot f^{s_{t-1}}(x_i) \\
&= \frac{1}{2^j} \sum_{X_i \in \mathcal{A}} \sum_{x_i \in \Delta} \Pr[X_i \in E_j^* \wedge X_i = x_i | s_{t-1}] \cdot f^{s_{t-1}}(x_i) \\
&= \frac{1}{2^j} \sum_{w \in \Omega(s_{t-1})} \Pr[w | s_{t-1}] \sum_{X_i \in E_j^*(w)} f^{s_{t-1}}(X_i(w)) \qquad (20)
\end{aligned}$$

The first equality is by definition of $F^{s_{t-1}}(\cdot)$ from (16). The second equality holds since the optimal policy must decide whether to schedule X_i (by time 2^j) without knowing the realization of X_i. Now for each $w \in \Omega(s_{t-1})$, due to submodularity of the function $f^{s_{t-1}}(\cdot)$, we get

$$\begin{aligned}
\sum_{X_i \in E_j^*(w)} f^{s_{t-1}}(X_i(w)) &\geq f^{s_{t-1}}(E_j^*(w)) \\
&= \sum_{i \in [m], f_i(s_{t-1}) < 1} \frac{f_i(E_j^*(w)) - f_i(s_{t-1})}{1 - f_i(s_{t-1})} \\
&\geq |C_j^*(w)| - |C(t, w)| \qquad (21)
\end{aligned}$$

Recall that $E_j^*(w)$ denotes the set of elements scheduled by time 2^j in OPT(conditional on w), as well as the realizations of these elements. The equality comes from the definition of $f^{s_{t-1}}$. The last inequality holds because $C(t, w) = \{i \in [m] \,:\, f_i(s_{t-1}) = 1\}$ and set $E_j^*(w)$ covers functions $C_j^*(w)$. Combining (20) and (21) gives:

$$\frac{1}{\ell_{\sigma(t)}} F^{s_{t-1}}(X_{\sigma(t)}) \geq \frac{1}{2^j} \Big(\mathbb{E}\left[|C_j^*| \;\middle|\; s_{t-1} \right] - \mathbb{E}\left[|C(t)| \;\middle|\; s_{t-1} \right] \Big). \qquad (22)$$

By deconditioning the above inequality (taking expectation over s_{t-1}) and using (17), we derive:

$$G_t \geq \frac{1}{2^j} \cdot \Big(\mathbb{E}[|C_j^*|] - \mathbb{E}[|C(t)|] \Big) \geq \frac{1}{2^j} \cdot \Big(\mathbb{E}[|C_j^*|] - \mathbb{E}[|C_j|] \Big),$$

where the last inequality uses $\mathbb{E}[C(t)]$ is non-decreasing and $t \leq 8\alpha 2^j$.

Now summing over all $t \in [8\alpha 2^{j-1}, 8\alpha 2^j)$ yields:

$$\Delta_j = \sum_{t=8\alpha 2^{j-1}}^{8\alpha 2^j} G_t \geq 4\alpha \Big(\mathbb{E}[|C_j^*|] - \mathbb{E}[|C_j|] \Big) = 4\alpha \Big(\mathbb{E}[|R_j|] - \mathbb{E}[|R_j^*|] \Big) \qquad (23)$$

Combining (23) and (19), we obtain:

$$4\alpha(\mathbb{E}[|R_j|] - \mathbb{E}[|R_j^*|]) \quad \leq \quad \alpha \mathbb{E}[|R_{j-1}|]$$

which simplifies to the desired inequality in Lemma 7.4. □

References

[1] Yossi Azar and Iftah Gamzu. Ranking with submodular valuations. In *SODA*, pages 1070–1079, 2011.

[2] Yossi Azar, Iftah Gamzu, and Xiaoxin Yin. Multiple intents re-ranking. In *STOC*, pages 669–678, 2009.

[3] Nikhil Bansal, Anupam Gupta, and Ravishankar Krishnaswamy. A constant factor approximation algorithm for generalized min-sum set cover. In *SODA*, pages 1539–1545, 2010.

[4] Amotz Bar-Noy, Mihir Bellare, Magnús M. Halldórsson, Hadas Shachnai, and Tami Tamir. On chromatic sums and distributed resource allocation. *Inf. Comput.*, 140(2):183–202, 1998.

[5] Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan. Detecting high log-densities: an $n^{1/4}$ approximation for densest k-subgraph. In *STOC*, pages 201–210, 2010.

[6] G. Calinescu and A. Zelikovsky. The polymatroid steiner problems. *Journal of Combinatorial Optimization*, 9(3):281–294, 2005.

[7] Robert D. Carr, Lisa Fleischer, Vitus J. Leung, and Cynthia A. Phillips. Strengthening integrality gaps for capacitated network design and covering problems. In *SODA*, pages 106–115, 2000.

[8] Deeparnab Chakrabarty and Chaitanya Swamy. Facility location with client latencies: Linear programming based techniques for minimum latency problems. In *IPCO*, pages 92–103, 2011.

[9] Moses Charikar, Chandra Chekuri, and Martin Pál. Sampling bounds for stochastic optimization. In *APPROX-RANDOM*, pages 257–269, 2005.

[10] K. Chaudhuri, B. Godfrey, S. Rao, and K. Talwar. Paths, trees, and minimum latency tours. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 36–45, 2003.

[11] Chandra Chekuri, Guy Even, and Guy Kortsarz. A greedy approximation algorithm for the group steiner problem. *Discrete Applied Mathematics*, 154(1):15–34, 2006.

[12] Chandra Chekuri and Martin Pál. A recursive greedy algorithm for walks in directed graphs. In *FOCS*, pages 245–253. IEEE Computer Society, 2005.

[13] Jittat Fakcharoenphol, Chris Harrelson, and Satish Rao. The k-traveling repairmen problem. *ACM Transactions on Algorithms*, 3(4), 2007.

[14] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004.

[15] Uriel Feige, László Lovász, and Prasad Tetali. Approximating min sum set cover. *Algorithmica*, 40(4):219–234, 2004.

[16] Naveen Garg, Goran Konjevod, and R. Ravi. A polylogarithmic approximation algorithm for the group steiner tree problem. *J. Algorithms*, 37(1):66–84, 2000.

[17] Michel X. Goemans and Jan Vondrák. Stochastic covering and adaptivity. In *LATIN*, pages 532–543, 2006.

[18] Daniel Golovin and Andreas Krause. Adaptive submodularity: A new approach to active learning and stochastic optimization. In *COLT*, pages 333–345, 2010.

[19] Anupam Gupta, Viswanath Nagarajan, and R. Ravi. Approximation algorithms for optimal decision trees and adaptive tsp problems. In *ICALP (1)*, pages 690–701, 2010.

[20] Anupam Gupta and Aravind Srinivasan. An improved approximation ratio for the covering steiner problem. *Theory of Computing*, 2(1):53–64, 2006.

[21] Eran Halperin and Robert Krauthgamer. Polylogarithmic inapproximability. In *Proceedings of the 35th Annual Symposium on Theory of Computing*, pages 585–594, 2003.

[22] Anton J. Kleywegt, Alexander Shapiro, and Tito Homem de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502, 2002.

[23] Goran Konjevod, R. Ravi, and Aravind Srinivasan. Approximation algorithms for the covering steiner problem. *Random Struct. Algorithms*, 20(3):465–482, 2002.

[24] Zhen Liu, Srinivasan Parthasarathy, Anand Ranganathan, and Hao Yang. Near-optimal algorithms for shared filter evaluation in data stream systems. In *SIGMOD Conference*, pages 133–146, 2008.

[25] David P. Williamson Martin Skutella. A note on the generalized min-sum set cover problem. *Operations Research Letters*, To appear, 2011.

[26] Kamesh Munagala, Utkarsh Srivastava, and Jennifer Widom. Optimization of continuous queries with shared expensive filters. In *PODS*, pages 215–224, 2007.

[27] Viswanath Nagarajan. *Approximation Algorithms for Sequencing Problems*. PhD thesis, Tepper School of Business, Carnegie Mellon University, 2009.

[28] Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*. Springer-Verlag, Berlin, 2003.